

Using Style to Understand Descriptions of Software Architecture

Gregory Abowd

Robert Allen

David Garlan

Computer Science Department Carnegie Mellon University Pittsburgh, PA 15213

Abstract

The software architecture of most systems is described informally and diagrammatically. In order for these descriptions to be meaningful at all, figures are understood by interpreting the boxes and lines in specific, conventionalized ways[5]. The imprecision of these interpretations has a number of limitations. In this paper we consider these conventionalized interpretations as architectural styles and provide a formal framework for their uniform definition. In addition to providing a template for precisely defining new architectural styles, this framework allows for the proof that the notational constraints on a style are sufficient to guarantee the meanings of all described systems and provides a unified semantic base through which different stylistic interpretations can be compared.

1 Introduction

Software architecture is an important level of description for software systems. At the software architectural level of abstraction, a system is typically described as a collection of interacting components. Components perform the primary computations of the system. Interactions between components include high level communication abstractions such as pipes, procedure calls, message passing, and event broadcast [8].

The software architecture of most systems is usually described informally and diagrammatically using boxes to represent components and lines to represent connections between components. In order for these descriptions to be meaningful at all, figures are typically understood by interpreting the boxes and lines in specific, conventionalized ways. For example, for one system boxes might represent filters and lines might represent pipes connecting ports of those filters. In another, boxes might represent abstract data types or objects, and lines might represent procedure calls.

While useful in documenting system designs, such diagrams — even with their conventional interpretations have a number of obvious limitations. Their imprecision makes it difficult to attach unambiguous meanings to the descriptions. It may be difficult to know when an implementation agrees with the more abstract description. It is virtually impossible to reason formally about the descriptions. It is difficult to compare two different descriptions, even for the same interpretation.

The most common solution to this problem is to constrain the architectural notation so that it maps directly into a well-defined execution model. For example, interfaces to components can be described solely in terms of their procedure signatures, and connectors can be restricted to procedure call. Other execution models include tasks with IPC and event-based systems [9]. When so constrained, descriptions can be mapped directly to module facilities of a programming language or other executable implementations, and can thereby be given precise meanings.

This approach, however, has a number of problems. Most significantly, it limits the expressiveness of architectural description to just those structures and building blocks supported by the target implementation language or system. If, for instance, architectural connections have to be phrased in terms of procedure calls, then higher-level interactions (such as protocols of communication) cannot be expressed directly. In addition, the relatively low level of description may make it difficult to reason about the architectural design.

We argue that what is needed instead is a way to give conventionalized interpretations of architectural descriptions a more formal basis. Designers can use the abstractions that are appropriate to the architectural description at hand, but still have the precision of a formal model. Our approach will be to view the collection of conventions that are used to interpret a description as defining an architectural style. We then show that architectural styles can be described formally in terms of a small set of semantic mappings, and illustrate how these mappings can be used to define formally two common architectural styles. The approach thus provides a framework in which new styles can be defined by a similar set of definitions. Finally, we demonstrate that having carried out this exercise, it is possible to use the formal descriptions to gain insight into the properties of a style and its relationships to other styles.

The main thrust of our argument and examples will be to demonstrate how to give meanings to architectural descriptions. In one respect this is nothing new; programming language researchers have been providing denotational semantics of programming languages for years. What is novel, however, is the specialization of the general semantic approach to the problem of understanding software architecture. As we will show, this can be done by providing a syntactic and semantic framework in which architectural styles

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

 ¹⁹⁹³ ACM 0-89791-625-5/93/0012...\$1.50

can be given meanings.

The specialization of general theory to this particular domain has a number of significant engineering benefits. First, it provides a template for formalizing new architectural styles in a uniform way, thereby simplifying and regularizing the way styles are given meanings. Second, it provides uniform criteria (in the form of proof obligations) for demonstrating that the notational constraints on a style are sufficient to provide meanings for all described systems. Third, it makes possible a unified semantic base through which different stylistic interpretations can be compared.

These benefits address a real problem for the growing number of industrial research and development efforts that are creating domain-specific architectural styles — or "reference architectures" — for specific product families [3, 4, 10]. To the extent that they formalize their architectural frameworks at all, the semantic descriptions produced by these efforts are typically developed from scratch, and each uses different, idiosyncratic conventions and semantic bases.¹ Semantic descriptions are therefore difficult to develop and, having developed them, few comparisons can be made between different development efforts.

To present our approach we begin in Section 2 by outlining the method we will use to define an architectural style as a mapping from syntactic descriptions to a (style-specific) semantic model. In Section 3, we formalize the syntactic domain as an abstraction of the box and line diagrams that are prevalent in current informal architectural descriptions. We then demonstrate for two particular architectural styles the definition of a semantic model to describe the overall behavior of a system and show how the architectural syntax can be mapped into that model by a formal style definition: we define a pipe-filter style in Section 4 and an implicit invocation, or event system, style in Section 5. Finally, in Section 6 we show how these semantic underpinnings support the analysis and comparison of styles.

We use the Z specification language to express the formal model in this paper [12]. Appendix A summaries the Z notation that we use in this paper.

2 What's in a Style?

In order to provide a precise meaning for architectural descriptions it is important to distinguish the abstract syntactic domain of architectural descriptions from the semantic domain of architectural meanings. Having done this we can then provide a map, or meaning function, from one to the other.

We take as our starting point the view that the syntactic domain of architectural description (among other things) supports the description of systems in terms of three basic syntactic classes: components, which are the locus of computation; connectors, which are the locus of communication; and configurations, which are collections of interacting components and connectors. Additionally, various style-specific concrete notations may be used to represent these visually, facilitate the description of legal computations and interactions, and constrain the set of describable systems.

A purely syntactic description may have some benefits as an informal design notation. For example, the connectors may be interpreted as defining data flows through the system. But as we argued in the introduction, such informal approaches have strong limitations. In particular, questions such as how components compute, what data is communicated, or how the flow of information is controlled, cannot be answered with any precision.

The purpose of this paper is to provide an improved basis for understanding the meaning of architectural descriptions. To do this we adopt the notion of architectural style as an interpretation from syntax to semantics, and outline a framework for precise style definition.

In this framework style definition starts with a formal definition of the syntactic domain in which architectures are described. In Section 3, we do this generically by providing formal definitions of components, connectors, and configurations. Next, for each style we must define a semantic model that captures both the static and dynamic meanings of the class of systems built in that style. Finally, as with a denotational approach to programming languages, we provide a mapping from the syntactic descriptions to the semantic model for the style. Given the nature of architectural descriptions, this amounts to the definition of three meaning functions that link the syntactic descriptions to their semantic counterparts. For style X, we would declare the meaning functions as partial functions from the abstract syntax to the semantic models.

 $\begin{array}{l} \mathcal{M}^X_{Comp}: \textit{Component} \rightarrow \textit{Comp}^X_{sem} \\ \mathcal{M}^X_{Conn}: \textit{Connector} \rightarrow \textit{Conn}^X_{sem} \\ \mathcal{M}^X_{Conf}: \textit{Configuration} \rightarrow \textit{Confisem} \end{array}$

Here Component is the abstract syntactic class of components (to be defined in Section 3) and $Comp_{sem}^X$ denotes the semantic model of a component in style X. Thus, \mathcal{M}_{Comp}^X is a meaning function from the general abstract syntax for components to the style-specific semantic model. Similar conventions are used for connectors and configurations.

The final step in the formal definition of an architectural style is to make explicit the constraints that this style imposes on the syntactic descriptions. Because the meaning functions are declared as partial functions on the syntactic domains, not every syntactic construct may have a meaning in a given style. Expressing these constraints explicitly carries a proof obligation to show that the meaning function is well-defined for all syntactic elements which meet the constraints. By making the constraints explicit we are precise about the descriptions that are reasonable in the style.

After we have formally defined an architectural style using the method outlined above, we have a foundation for further analysis of the style. We discuss two different forms of analysis in this paper. The first form of analysis is within a particular style, identifying important substyles that can be understood as further syntactic restrictions on a more general style. The second form of analysis is between styles, which we exemplify by comparing different semantic models to see if they share similar properties (in this case we check whether configurations can be collapsed and represented as components).

To summarize, the steps we will follow are:

- formalize abstract syntax for architectures
- for a given style:
 - define the semantic model
 - discuss concrete syntax for the style
 - define the mapping from abstract syntax into semantic model

¹The alternative, attempted by some, is to legislate that different groups use a given, prescribed procrustean style of architectural description. As noted above, this has serious limitations.

- make explicit the constraints on the syntax
- demonstrate analysis within and between formally defined architectural styles.

3 The Abstract Syntax of Software Architectures

From an abstract, generic point of view the basic syntactic elements of an architectural description are components, connectors, and configurations of components and connectors.

3.1 Components

Components are the active, computational entities of an architecture. They accomplish tasks through internal computation and external communication with the rest of the system. The relationship between a component and its environment is defined explicitly as a collection of interaction points, or ports. We can also differentiate between components with the same port interface based on a description of the computation they perform. At the abstract level of a component, we model this reference to computational behavior with a placeholder for some concrete syntactic description.

For the moment, we are not concerned with details of the construction of ports or the computational description for components, so we model these as given sets. An architectural component, as a syntactic entity, is modeled as a collection of ports together with a description of its computation. We use the Z schema construct to define the new type *Component* to represent the syntax of components in our architectural specification.

[PORT, COMPDESC]

Component	
ports : P PORT	
description : COMPDESC	

3.2 Connectors

Connectors define the communication between components. Each connector provides a way for a collection of ports to come into contact. A connector, rather than being bound unchangeably to specific ports on specific components, provides placeholders for these ports, as *roles* in the communication. The description of the precise communications protocol provided within a connector is separated in the same way that we separate the computation description in a component from its port interface. The exact language used to describe this communication behavior is an issue for the connector's concrete syntax, and we represent it here as only a placeholder.

Again, we are not yet concerned with the details of roles or communication description, so we introduce them as given sets in this specification. An architectural connector is modeled as a collection of roles together with a description of its communication protocol, as defined in the schema Connector.

[ROLE, CONNDESC]

_ Connector _	
roles : P RO.	LE
description :	CONNDESC

3.3 Configurations

Named instances of components and connectors are combined to form configurations. Names for instances of components and connectors are taken from two more given sets. The names used to identify component and connector instances are also used to identify instances of ports and roles. We introduce type synonyms for named ports and roles.

[COMPNAME, CONNNAME] PortInst == COMPNAME × PORT RoleInst == CONNNAME × ROLE

We use a partial function, components, to represent naming the set of components in a particular configuration in the schema Configuration below. Similarly, we use the partial function connectors for naming connectors in a configuration. The interfaces of components and connectors are attached to reflect their composition. Roles are filled by particular ports, and communication occurs along the connections. This is modeled as a partial function, attachment, from role instances to port instances, reflecting our intuition that roles represent placeholders for particular ports. While a port may fill many roles, meeting the needs of several different communications, a role may have at most one port that fills it.

In addition to declaring the attributes for the named components and connectors and the attachment between roles and ports, the schema *Configuration* includes two additional constraints (below the separating line) that must be satisfied by all configurations. The first constraint is a predicate that ensures that any role instance in the *attachment* is a role for some named connector in the configuration. The second constraint ensures a similar fact for port instances and the named components. Together, these two constraints enforce a lexical scoping on attachments within a configuration.

Configuration
components : COMPNAME → Component
connectors : $OONNNAME \rightarrow Connector$
$attachment: RoleInst \rightarrow PortInst$
$\forall cn : CONNNAME; r : ROLE \\ (cn, r) \in dom attachment$
• $cn \in dom \ connectors \land r \in (connectors(cn)).roles$
∀ cn : COMPNAME; p : PORT
$ (cn, p) \in ran attachment$
• $cn \in dom \ components \land p \in (components(cn)).$ ports

4 The Pipe-Filter Style

In this section, we show how the framework can be used to define the pipe-filter Style (PF). This style is representative of coarse-grained dataflow systems, such as those supported by Unix pipes.

4.1 Semantic Model

The first part of defining a style is to provide a semantic model for the components, connectors, and configurations of the style. In general, this is perhaps the hardest part of the process, since to do this properly we must come to grips with the intuition behind the use of the style. In the case of PF, an appropriate formal description of the semantic domain already exists [1, 2]. Here we will use only those aspects of the model that are necessary to illustrate the basic ideas.

The PF style interprets components as filters, which are typed stream transducers. These can be modeled as state machines that receive their input and place their output as sequences on data ports. We do not wish to uncover the details of how the internal state and data are described, so we declare them as given sets in our specification. Data ports define the interfaces for filters and we also introduce them as a given set in our model. These are to be distinguished from the ports that form the interface for unnamed components in the syntactic descriptions.

[STATE, DATA, DATAPORT]

In order to define the behavior of a filter, we must know its input and output ports and the type of data that may be passed along each port. This latter information can be represented by a (partial) function from data ports to their alphabet. At any point in time, the ports of the filter will hold all data (as a sequence) that has been received (for input ports) or produced (for output ports) but not yet removed. The state machine behavior of the filter is modeled as a transition function that takes an internal state and data configuration and results in a new internal state and data configuration. In addition we can identify a starting internal state. This information about a filter is summarized in the schema *Filter*. Some constraints on filters that we enforce are:

- input and output data ports are distinct (first predicate below);
- a filter transition is determined by looking at data on the input ports only and results in information provided to the output ports only (the final predicate below).

Filter
inputs, outputs : P DATAPORT
$alphabet : DATAPORT \rightarrow P DATA$
states : P STATE
start : STATE
$\begin{array}{l} \textit{transitions} : (STATE \times (DATAPORT \dashrightarrow seq DATA)) \\ \leftrightarrow (STATE \times (DATAPORT \dashrightarrow seq DATA)) \end{array}$
$inputs \cap outputs = \emptyset$ dom alphabet = inputs \cup outputs start \in states
$\forall s_1, s_2 : STATE; ps_1, ps_2 : DATAPORT \rightarrow seq DATA$
• $((s_1, ps_1), (s_2, ps_2)) \in transitions \Rightarrow$
$s_1 \in states \land s_2 \in states$
$\wedge \text{dom } ps_1 = inputs \land \text{dom } ps_2 = outputs$
$\land (\forall i: inputs \bullet ran(ps_1(i)) \subseteq alphabet(i))$
$\land (\forall o: outputs \bullet ran(ps_2(o)) \subseteq alphabet(o))$

We can define an operational semantics for the computational behavior of a filter. At any point time, a filter is defined by its current internal state, constrained to be in the set of possible states for the filter, and the data at each of its input and output ports (which must be in the alphabet of that port).

FilterState f:Filter curstate:STATE instate.outstate:DATAPORT
$curstate \in f.states$
dom $instate = f.inputs$
dom $outstate = f.outputs$
$\forall p: f.inputs \bullet ran(instate(p)) \subseteq f.alphabet(p)$
$\forall p: f.outputs \bullet ran(outstate(p)) \subseteq f.alphabet(p)$

A single computation for a filter transforms some input data into output data. The order of data is preserved, so input data is consumed in the order it arrived and output data is provided in the order it is produced. The result of a computation step for a filter is the removal of some data off the input ports, a transformation of that data, which will depend on the filter's current state, a change in the current state and the addition some data to the output ports. The schema *FilterCompute* encapsulates just such a computational step. We make use of the Δ convention to describe this transition from one state of the filter to another (see Appendix A).

FilterStep	_
Δ FilterState	
f' = f	
$\exists in, out : DATAPORT \rightarrow seq DATA$ • ((curstate, in), (curstate', out)) $\in f.transitions$ $\land \forall p : f.inputs$ • instate(p) = indata(p) ^ instate'(p) $\land \forall p : f.outputs$ • outstate'(p) = outstate(p) ^ outdata(p)	

The data ports of transducers are connected by pipes, which we model as typed streams of data. Each pipe has a distinct source and sink for receiving and sending data.

<i>Pipe</i>	
source, sink : DATAPORT alphabet : P DATA	
source \neq sink	

The protocol or behavior of a pipe is defined by giving its transmission policy. At any point in time, the pipe has some data residing at its source ports and some data at its sink ports.

_ PipeState	
p: Pipe	
sourcedata : seq DATA	
sinkdata : seq DATA	
ran sourcedata $\subseteq p.alphabet$	
$\operatorname{ran} sinkdata \subseteq p.alphabet$	

A single step in the behavior of a pipe results in some nonempty subsequence of data being removed from the source data ports, in the order in which it arrived there, and being delivered, unchanged in content and order, to the sink data ports.

$$\begin{array}{c} PipeStep \\ \underline{\Delta PipeState} \\ \hline p = p' \\ \exists deliver : seq DATA \\ | # deliver > 0 \\ \bullet \quad deliver \widehat{} sourcedata' = sourcedata \\ \land \quad sinkdata' = sinkdata \widehat{} deliver \end{array}$$

A configuration is then modeled as a set of filters connected by pipes. We disallow name clashes between the data ports of distinct filters and pipes. The interaction is modeled by identifying each pipe *source* with a unique filter output and each pipe *sink* with a unique filter input.

InteractingFilterSet_ filters : P Filter pipes : P Pipe $\forall f_1, f_2 : filters$ $f_1 \neq f_2$ • $(f_1.inputs \cup f_1.outputs) \cap (f_2.inputs \cup f_2.outputs) = \emptyset$ $\forall p_1, p_2 : pipes$ $p_1 \neq p_2$ • { p_1 .source, p_1 .sink} \cap { p_2 .source, p_2 .sink} = \emptyset $\forall p : pipes$ • $\exists f_1, f_2$: filters $p.source \in f_1.outputs$ $p.sink \in f_2.inputs$ Λ $f_1.alphabet(p.source) = p.alphabet$ ٨ $f_2.alphabet(p.sink) = p.alphabet$ Λ

The behavior of an interacting set of filters is defined in terms of the behaviors of the constituent filters and pipes. One step in this behavior is either a computation step for one filter or a transmission step for one pipe, all else remaining unchanged. Details of this behavioral specification have been omitted here but can be found in [2].

4.2 Concrete Syntax

The second part of a style definition is the creation of a stylespecific concrete syntax. While the details of such syntax are important, in this paper we are more concerned with understanding the relationship between these descriptions and their associated meanings. In that regard, it is enough to know that there exist filter and pipe description languages that determine the interesting subset of the possible component and connector descriptions in the PF style. Formally, we represent these languages as subsets of the respective description languages introduced in Section 3.

FilterDescriptions : P COMPDESC PipeDescriptions : P CONNDESC

For concreteness, Figure 1 illustrates the definition of a filter that capitalizes its character input stream using one notation developed for this style [2].

4.3 Meaning Functions

The third part of a style description is to define the meaning of the syntactic constructs in terms of the semantic model.

As indicated in Section 2 to give meaning to components, we need to specify a partial function of the form:

```
inputs: char in;
outputs: char out;
execution:
    char c;
    while (TRUE) {
        c = read(in);
        if (c >= 'a' && c <= 'z') {write(out,c+'A'-'a');}
        else {write(out,c);}
    }
}
```

Figure 1: Concrete Description of a Capitalizing Filter

 $\mathcal{M}_{Comp}^{X}: Component \rightarrow Comp_{sem}^{X}$

From the definition of *Filter*, we can see that it is possible for two filters to be identical up to naming of data ports and states. Therefore, we can define an equivalence relation on elements in *Filter*. We treat two filters as equivalent if and only if there is an isomorphism between their states, and their input and output data ports that preserves the behavior defined by their transition functions. This equivalence relation is denoted by \equiv_{fil} . The detailed definition of \equiv_{fil} is not given below, though it is straightforward.

 $= \equiv_{fil} = : Filter \leftrightarrow Filter$

The meaning function for PF components, written below as \mathcal{M}_{Comp}^{PF} , identifies the syntactic element Component with an equivalence class of filters. So in this example, $Comp_{sem}^X$ is replaced by sets of filters, or **P** Filter. In order to complete the mapping from syntax to semantics, we need to have an injective function, called DataPort below, from named instances of the syntactic ports to the semantic data ports. Among other things, DataPort will help to distinguish between computationally equivalent filters.

$$\begin{array}{l} DataPort: PortInst \rightarrow DATAPORT\\ \mathcal{M}_{Comp}^{PF}: Component \rightarrow \mathbf{P} \ Filter\\ \forall c: Component; f_1, f_2: Filter\\ \mid f_1 \in \mathcal{M}_{Comp}^{PF}(c)\\ \bullet \ f_2 \in \mathcal{M}_{Comp}^{PF}(c) \Leftrightarrow f_1 \equiv_{fil} f_2\\ \forall c: Component; n: COMPNAME\\ \mid c \in \text{dom} \ \mathcal{M}_{Comp}^{PF}\\ \bullet \ \exists f: \mathcal{M}_{Comp}^{PF}(c)\\ \bullet \ DataPort(\{n\} \times c.ports\}) = (f.inputs \cup f.outputs)\\ \end{array}$$

In Section 4.4 we will discuss what constraints on components must hold in order to give them meaning in the PF style. That is, we will explicitly define the domain of the function \mathcal{M}_{Comp}^{PF} .

Connectors are given meaning in PF by interpreting them as pipes. The concrete syntax for pipes specifies the type of data transmitted. Two pipes are considered equivalent if they have the same alphabets. Of course, in the context of a set of interacting filters, the pipes are distinguished by the dataports they connect.

$$\begin{array}{c} \mathcal{M}_{Conn}^{PF}: Connector \leftrightarrow \mathbf{P} \ Pipe \\ \hline \forall \ c: \ Connector; \ p_1, p_2: \ Pipe \\ \mid p_1 \in \mathcal{M}_{Conn}^{PF}(c) \\ \bullet \ p_2 \in \mathcal{M}_{Conn}^{PF}(c) \Leftrightarrow p_1. alphabet = p_2. alphabet \\ \end{array}$$

We can now define the meaning of configurations in the PF style. Named components are interpreted as filters and connectors are realized as pipes. The attachments are realized by equating pipe sources with unique filter data ports and pipe sinks with unique filter input data ports. To do this we select appropriate elements from the classes defined by the meaning functions \mathcal{M}_{Comp}^{PF} and \mathcal{M}_{Conn}^{PF} . In the syntactic domain, we declare that source and sink are distinct roles for connectors.

source, sink : ROLE $source \neq sink$ $\mathcal{M}_{Conf}^{PF} : Configuration \rightarrow InteractingFilterSet$ $\forall cfg : dom \mathcal{M}_{Conf}^{PF} \bullet$ $(\mathcal{M}_{Conf}^{PF}(cfg)).filters =$ $\{n : COMPNAME; c : Component; f : Filter$ $\mid (n, c) \in cfg.components$ $\wedge f \in \mathcal{M}_{Comp}^{PF}(c)$ $\wedge f.outputs \cup f.inputs = DataPort(\{n\} \times c.ports\})$ $\bullet f\}$ \wedge $(\mathcal{M}_{Conf}^{PF}(cfg)).pipes =$ $\{n : CONNNAME; c : Connector; p : Pipe$ $\mid (n, c) \in cfg.connectors$ $\wedge p \in \mathcal{M}_{Conn}^{PF}(c)$ $\wedge p.source = DataPort(cfg.attachment(n, source))$ $\wedge p.sink = DataPort(cfg.attachment(n, sink))$ $\bullet p\}$

4.4 Syntactic Constraints

The final part of defining a style is to make explicit the syntactic preconditions that must be satisfied in order to translate to the semantic domain. Since the meaning functions are partial, only a subset of all components, connectors and configurations are given a meaning in the PF style. This corresponds to the intuition that only some architectural descriptions represent valid pipe-filter systems. In particular, for components we demand that the computation associated with the component can be defined using the concrete language of FilterDescription and that the named component ports can be realized as data ports of some filter. We can express these syntactic constraints in Z by use of schema inclusion in which the original specification of type Component is included in the specification of syntactically legal PF components and then further constrained. (See Appendix A for further details on schema inclusion.)

LegalPFComponen Component	
description ∈ Filte	rDescriptions

By specifying this explicit syntactic constraint, we are actually asserting two things. First, only component descriptions that satisfy this constraint can be legally interpreted as a filter. This is equivalent to asserting that the domain of \mathcal{M}_{Comp}^{PF} is LegalPFComponent.

dom $\mathcal{M}_{Comp}^{PF} = Legal PF Component$

Second, this assertion results in a proof obligation that we have not invalidated our definition of \mathcal{M}_{Comp}^{PF} . In other

words, we must prove that given any legal PF component, we can apply \mathcal{M}_{Comp}^{PF} to obtain a filter. We must show that

$$\forall c: LegalPFComponent \bullet \mathcal{M}_{Comp}^{PF}(c) \neq \emptyset$$

This amounts to demonstrating that

or that the function DataPort is reasonably constructed. Therefore, the domain restriction to \mathcal{M}_{Comp}^{PF} is valid.

Similarly, we constrain the definition of connectors to be those having a concrete description interpretable as a stream alphabet and having only two roles, *source* and *sink*.

LegalPFConnector Connector	
$description \in PipeDescriptions \\ roles = \{source, sink\}$	

Once again, we formally restrict the meaning function to cover legal values.

dom $\mathcal{M}_{Conn}^{PF} = Legal PFC onnector$

This also results in a proof obligation. Since \mathcal{M}_{Conn}^{PF} as defined could be total, however, the proof is trivial.

As one might expect, the constraints we enforce on configurations are more complex. For the pipe and filter style defined above these are:

- 1. Each named component is a legal filter.
- 2. Each named connector is a legal pipe.
- 3. Every pipe source is attached to a unique filter output with the same alphabet.
- 4. Every pipe sink is attached to a unique filter input with the same alphabet.

In the following schema, the first two predicates below the line express the first two constraints above. The third predicate below states that all pipe sources and sinks are attached to some named ports. The fourth predicate says that the attachment function is injective, that is, no two sources or sinks can be attached to the same port instances. The last two predicates express the alphabet constraint.

A straightforward argument shows that any syntactically legal configuration can be assigned a meaning by \mathcal{M}_{Conf}^{PF} , so we restrict its domain to LegalPFConfig.

dom $\mathcal{M}_{Conf}^{PF} = Legal PFC on fig$

This concludes our formal definition of the PF style. In Section 6 we investigate other syntactic constraints that can be used to define PF substyles and discuss some analysis that can be performed on the semantic domain of PF.

5 Event System Style

In this section, we show (more briefly) how the same method of definition for the PF style can be used to describe another common architectural style, the event system (ES). Event systems are increasingly important as a flexible tool integration technique, since they allow the implicit invocation of tools when some other tool announces an event[6, 11].

For the purposes of this paper we will treat each component in an event system as a collection of methods sharing a state. A component responds to an incoming method by transforming its internal state and announcing some events. Connection in the system consists of an association between events and the methods that should be invoked when those events are announced.

5.1 Semantic Domain

The ES style interprets components as *objects* with a vocabulary of methods and events. Here we will model an object as a state machine with a transition function relating method invocations to state transitions and event announcement.

[METHOD, EVENT]

Object
methods : P METHOD
events : P EVENT
states : P STATE
start : STATE
$transitions: (METHOD \times STATE)$
$\rightarrow (STATE \times \mathbf{P} EVENT)$
$start \in states$
dom transitions = methods \times states
ran transitions \subseteq { s : states; es : $P EVENT \bullet (s, es)$ }

The ES style interprets connectors as distributors, which take announced events and transform them into method invocations. Our model of a distributor below is understood as saying that whenever any event in *events* is announced, then every method in *methods* must be invoked.

 Distributor	
events : P EVENT	
methods : P METHOD	

A collection of objects and distributors are joined to form a set of interacting objects. Unlike the PF example, the semantic model does not constrain the topology of the connectors. Indeed, the only constraint we express in the semantic domain is that there be no name clash between the events and methods of the objects. (That is, in this model events and methods are uniquely associated with specific objects.)

InteractingObjectSet	
objects : P Object	
distributors : P Distributor	
$\forall o_1, o_2 : objects \bullet$	
$(o_1.events \cap o_2.events \neq \emptyset \lor$	
$o_1.methods \cap o_2.methods \neq \emptyset) \Rightarrow o_1 = o_2$	

We have defined the static view of the semantic model for event systems. We will only outline the model for the dynamic behavior of ES. At any point in time, each object in the system will be in some legal state and will have some methods that have been invoked but not executed. The system will also hold a set of announced events that have not yet been interpreted and distributed as method invocations to the relevant objects.

A change in the system results when either a single object performs one of its pending invoked methods or when an announced event is distributed as method invocations to the relevant objects.

5.2 Concrete Syntax

A concrete syntax for events systems can be developed as an extension of regular programming languages [13]. The

```
for Package_1
    declare Event_1 X : Integer,
    declare Event_2
    when Event_3 => Method_1 B
end for Package_1
for Package_2
    declare Event_3 A,B : Integer;
    when Event_1 => Method_2 X
    when Event_2 => Method_4
end for Package_2
```

Figure 2: Event System Description	Example	e
------------------------------------	---------	---

details of these extensions are not particularly important for this discussion. These concrete descriptions define a subset of allowable computation and communication descriptions.

ObjectDescriptions : P COMPDESC DistributorDescriptions : P CONNDESC

For example, Figure 2 illustrates a concrete syntax for the communication description extension that allows an Ada package interface to specify events announced by that package and the method to be invoked when an event is announced by some other package [7].

5.3 Meaning Functions

The definition of meaning functions for ES proceeds exactly as for PF. The meaning function for ES components, written \mathcal{M}_{Comp}^{ES} , associates the syntactic elements of *Component* with equivalence classes of objects. Equivalence between objects is denoted by \equiv_{obj} .

As with PF, in order to complete the mapping from syntax to semantics we need to identify ports and roles (the syntactic elements) with events and methods (the semantic elements). Named port instances are identified as either a method or event, but not both. Roles are identified as either event roles or method roles.

```
\begin{array}{l} EventasPort: PortInst \longrightarrow EVENT\\ MethodasPort: PortInst \longrightarrow METHOD\\ EventRoles: $$$$$$$$$$$$$$$ ROLE\\ MethodRoles: $$$$$$$$$$$$$$$$$$$$$ ROLE\\ \end{array}
```

(dom EventasPort, dom MethodasPort) partition PortInst

 $\forall n, n' : COMPNAME; p : PORT$ • $(n, p) \in \text{dom EventasPort}$ $\Leftrightarrow (n', p) \in \text{dom EventasPort}$ $\land (n, p) \in \text{dom MethodasPort}$ $\Leftrightarrow (n', p) \in \text{dom MethodasPort}$

```
(EventRoles, MethodRoles) partition ROLE
```

The ES style interprets components as objects, matching the methods and events of the object to corresponding ports.

$$\mathcal{M}_{Comp}^{ES}: Component \rightarrowtail \mathsf{P} \ Object$$

$$\forall c: Component; o_1, o_2: Object$$

$$\mid o_1 \in \mathcal{M}_{Comp}^{ES}(c)$$

$$\bullet o_2 \in \mathcal{M}_{Comp}^{ES}(c) \Leftrightarrow o_1 \equiv_{obj} o_2$$

$$\forall n: COMPNAME; c: dom \mathcal{M}_{Comp}^{ES}$$

$$\bullet \exists o: \mathcal{M}_{Comp}^{ES}(c)$$

$$\bullet EventasPort^{\circ}(o.events) \cup MethodasPort^{\circ}(o.methods)$$

$$= \{n\} \times c.ports$$

The ES style interprets connectors as distributors. The distributor represented must have the same number of events and methods as the connector. Note that we are essentially defining the criteria for equivalence of distributors.

$$\mathcal{M}_{Conn}^{ES}: Connector \rightarrowtail \mathbf{P} \ Distributor$$

$$\forall c: Connector; d: Distributor$$

$$| d \in \mathcal{M}_{Conn}^{ES}(c)$$

$$\bullet \qquad \#(d.events) = \#(c.roles \cap EventRoles)$$

$$\land \qquad \#(d.methods) = \#(c.roles \cap MethodRoles)$$

The meaning of a configuration is derived from the meaning of its components, its connectors, and the attachment function. The attachment links events announced by an object to the same event received by one or more distributors. Also the attachment links methods received by an object to the same method invoked by one or more distributors.

$$\begin{split} \mathcal{M}^{ES}_{Conf}: Configuration & \leftrightarrow InteractingObjectSet \\ \forall cfg: dom \mathcal{M}^{ES}_{Conf} \bullet \\ & (\mathcal{M}^{ES}_{Conf}(cfg)).objects = \\ & \{n: dom cfg.components; c: Component; o: Object \\ & | cfg.components(n) = c \\ & \land o \in \mathcal{M}^{ES}_{Comp}(c) \\ & \land EventasPort^{\sim}(o.events) \cup MethodasPort^{\sim}(o.methods) \\ & = \{n\} \times c.ports \\ & \bullet o\} \\ & \land \\ & (\mathcal{M}^{ES}_{Conf}(cfg)).distributors = \\ & \{n: dom cfg.connectors; c: Connector; d: Distributor \\ & | cfg.connectors(n) = c \\ & \land d \in \mathcal{M}^{ES}_{Conn}(c) \\ & \land (\forall r: c.roles; (n', p) \in dom EventasPort \\ & \bullet cfg.attachment(n, r) = (n', p) \Leftrightarrow \\ & EventasPort(n', p) \in d.events) \\ & \land (\forall r: c.roles; (n', p) \in dom MethodasPort \\ & \bullet cfg.attachment(n, r) = (n', p) \Leftrightarrow \\ & MethodasPort(n', p) \in d.methods) \\ & \bullet d \} \end{split}$$

5.4 Syntactic Constraints

The syntactic constraints in the ES style can be expressed by making explicit the domain for the meaning functions. For components, we simply restrict interpretation to those whose computation can be described using the concrete language of *ObjectDescriptions*.

LegalObject _____ Component description ∈ ObjectDescriptions Similarly for distributors, we restrict the abstract syntax to include only those connectors whose protocol can be described by the language of *DistributorDescriptions*.

LegalDistri	butor
Connector	
description	$\in Distributor Descriptions$

A legal configuration is one in which the components are legal objects, the connectors are legal distributors, and attachments only occur between event roles and event ports or between method roles and method ports.

LegalESConfig Configuration	•
$ \forall c : ran \ components \bullet c \in LegalObject \\ \forall c : ran \ connectors \bullet c \in LegalDistributor $	
$\forall n : CONNNAME; m : COMPNAME;role : ROLE; port : PORT ((n, role), (m, port)) \in attachment • role \in EventRoles \Leftrightarrow (m, port) \in dom EventasPort$	

The domains of the meaning functions are accordingly defined.

 $\begin{array}{l} \operatorname{dom} \mathcal{M}_{Comp}^{ES} = LegalObject \\ \operatorname{dom} \mathcal{M}_{Conn}^{ES} = LegalDistributor \\ \operatorname{dom} \mathcal{M}_{Conf}^{ES} = LegalESConfig \end{array}$

6 Analysis Using Architectural Style

One of the main reasons to formalize architectural style is to gain analytic leverage. In this section we present two examples of the kind of analysis of an architectural style that is possible within our formal framework.

6.1 Defining Architectural Substyles

It is common for one style to be understood in terms of another. Many of these *substyles* can be understood as additional constraints on the syntax of the more general style. For example, in the PF style we can identify the following common substyles:

- disallowing feedback loops, or cycles;
- restriction to a pipeline; and
- allowing only a fan-out of components.

The nature of pipes permits us to consider the topology of a PF configuration as a directed graph. We can derive the connection between two components by determining if any of their ports are attached to a common pipe.

```
\begin{array}{l} PFGraph \\ LegalPFConfig \\ connect : COMPNAME \leftrightarrow COMPNAME \\ \hline \\ \hline \\ connect = \\ \{c_1, c_2: dom \ components; \ pipe : dom \ connectors \\ | \ attachment(pipe, \ source) = (c_1, p_1) \\ \land attachment(pipe, \ sink) = (c_2, p_2) \\ \bullet (c_1, c_2) \} \end{array}
```

A PF system with no feedback loops is one in which the connection graph is acyclic.

AcyclicPF_	·····		
PFGraph			
id COMPN	$ME \cap connect^+$	- Ø	
	In D . Connect		

To express acyclic pipe-filter architectures as an independent style, we restrict the meaning function \mathcal{M}_{Conf}^{PF} to configurations satisfying *Acyclic*. The other meaning functions are the same as for the general PF style.

$$\begin{split} \mathcal{M}^{Acyclic}_{Comp} &: Component \to \mathbf{P} \ Filter \\ \mathcal{M}^{Acyclic}_{Conn} &: Connector \to \mathbf{P} \ Pipe \\ \mathcal{M}^{Acyclic}_{Conf} &: Configuration \to Interacting Filter Set \\ \hline \mathcal{M}^{Acyclic}_{Comp} &= \mathcal{M}^{PF}_{Comp} \\ \mathcal{M}^{Acyclic}_{Conn} &= \mathcal{M}^{PF}_{Conn} \\ \mathcal{M}^{Acyclic}_{Conn} &= \{Acyclic \bullet \theta \ Configuration\} \triangleleft \mathcal{M}^{PF}_{Conf} \end{split}$$

Restriction to a pipeline means that we can view the connection graph as a sequence of components, with each component in the pipeline sequence connected to the component after it in the pipeline.

Pipeline _		
PFGraph		
∃ filters : s	seq COMPNAME	
ran filt	$ters = dom \ components$	
• conne	$xt = \{i: 1 \dots (\# filters - 1)\}$	
	• $(filters(i), filters(i+1))$	

A PF substyle allowing only fan-out has a connection graph whose inverse is a function, that is, components are connected to a unique parent component that provides its input.

$connect \sim \in$	COMPNAME -++ COMPNAME	
PFG raph		
FanOut		_

Garlan and Notkin have used the event system model to investigate the differences between various implementations of an implicit invocation mechanism [6]. Their examples concentrate on restrictions to the kinds of events that objects can announce and the form of the event to method binding that a distributor allows. Since we have left the interpretation of events and methods open and allow distributors to bind events to methods arbitrarily, all of those styles are substyles of ES as it appears in this paper.

Another substyle of ES is one with a global event name space. In this substyle, different objects can announce the same event. In that sense, it is the component port that uniquely identifies events and not the named port instance. This substyle can be expressed as a constraint that events must be distributed to the same method invocations regardless of which component announces the event. One way to express this constraint syntactically is shown below.

6.2 Relating Semantic Domains

One desirable property of an architectural description is hierarchy. In a hierarchical description components or connectors may themselves be represented as a configuration. For example, in the pipe and filter style, we might want to allow one filter to be expandable into a configuration of pipes and filters. By defining a style formally, we can understand what properties of the semantic domain might make this kind of description meaningful.

For example, Allen and Garlan showed formally that in the pipe and filter style it is semantically meaningful to decompose a component (filter) into a configuration of pipes and filters [2]. In their treatment, the decomposition is meaningful when the behavior of the unbound ports of the associated configuration matches the behavior of ports of an equivalent filter. In brief, the proof consists of the construction of a relation between a set of interacting filters and a single filter.

$_collapse_{PF}_: InteractingFilterSet \leftrightarrow Filter$

This result means that we can now expand the concrete description language of filters to include hierarchical decomposition without altering our useful and intuitive understanding of the PF style.

This result leads us to ask whether a similar result holds for any other style. For example, in the event system style, does there exist a similar relation between collections of interacting objects and single objects? In other words, does there exist a relation

$| _collapse_{ES}_: InteractingObjectSet \leftrightarrow Object$

such that the external method/event behavior of the set of objects matches that of the collapsed object?

Without going into too much detail, we can see that in general this result does not hold for event systems. When a set of interacting objects is collapsed into a single object any event-method connections internal to the set of objects will result in a computation that cannot be made to correspond to any visible method invocation.

This is a useful result, because it tells us that if we want to provide hierarchical event systems we must do one of two things. Either we have to change the semantic model, or we have to find ways to restrict the class of descriptions to a subset that allows hierarchical decomposition. In the former case we would need to view method invocation as non-atomic. In the latter case we might restrict decompositions to be configurations that do not have any internal event-method bindings.

7 Conclusion

We have argued that a formal approach to architectural style permits the precise interpretation and analysis of architectural descriptions. This has two important benefits. First, precision facilitates effective communication about systems at the architectural level. Misunderstandings inherent in ambiguous specifications can be avoided without abandoning the architectural paradigm. Second, a formal understanding of classes of systems permits the development of specialized analysis techniques as well as comparison between styles.

In addition to these immediate benefits, a precise understanding of style represents a necessary first step toward automated support for software architectural design and development. Through an understanding of both the structural constraints and the semantic underpinnings of architectures, tools and environments can be developed that effectively support the design process. As a first step in this direction, we have developed a software environment framework based on this model for style definition. The common elements of component, connector, configuration, and hierarchy are directly supported by the environment, while its open structure supports the development and integration of tools that take advantage of style-specific structural and semantic properties. Because of the generality of structured style definition, tools developed for one style may be reused for any style that has certain properties in common with the original.

Acknowledgments

The authors would like to thank various colleagues whose comments on this work have helped us to clarify our thoughts, especially Dave Wile, Marc Graham, Mary Shaw, Jeannette Wing, Daniel Jackson, John Ockerbloom and Amy Moormann Zaremski. This research was sponsored by the National Science Foundation under Grant Number CCR-9112880 and by Siemens Corporate Research. The views and conclusions contained in this document are those of the authors and should not be interpreted as representing the official policies, either expressed or implied, of the U.S. Government or of Siemens Corporation.

References

- ALLEN, R., AND GARLAN, D. A formal approach to software architectures. In *Proceedings of IFIP'92* (September 1992), J. van Leeuwen, Ed., Elsevier Science Publishers B.V.
- [2] ALLEN, R., AND GARLAN, D. Towards formalized software architectures. Tech. Rep. CMU-CS-92-163, Carnegie Mellon University, School of Computer Science, July 1992.
- [3] Proceedings of the Workshop on Domain-Specific Software Architectures (Hidden Valley, PA, July 1990), Software Engineering Institute.
- [4] EARL, A. A reference model for computer assisted software engineering environment frameworks. Tech. Rep. HPL-SEG-TN-90-11, Hewlett Packard Laboratories, Bristol, England, August 1990.
- [5] FREEMAN, P., AND A.I.WASSERMAN. Tutorial on Software Design Techniques. IEEE Computer Society Press, 1976.

- [6] GARLAN, D., AND NOTKIN, D. Formalizing design spaces: Implicit invocation mechanisms. In VDM'91: Formal Software Development Methods (Noordwijkerhout, The Netherlands, October 1991), Springer-Verlag, LNCS 551, pp. 31-44.
- [7] GARLAN, D., AND SCOTT, C. Adding implicit invocation to traditional programming languages. In Proceedings of the Fifteenth International Conference on Software Engineering (Baltimore, MD, May 1993).
- [8] GARLAN, D., AND SHAW, M. An introduction to software architecture. In Advances in Software Engineering and Knowledge Engineering, Volume I (New Jersey, 1993), V. Ambriola and G. Tortora, Eds., World Scientific Publishing Company.
- [9] LUCKHAM, D. C., AND VERA, J. Event-based concepts and language for system architecture. Working draft, October 1992.
- [10] METTALA, E., AND GRAHAM, M. H. The domainspecific software architecture program. Tech. Rep. CMU/SEI-92-SR-9, Carnegie Mellon Software Engineering Institute, June 1992.
- [11] REISS, S. Connecting tools using message passing in the Field Environment. *IEEE Software* 7, 4 (July 1990), 57-66.
- [12] SPIVEY, J. The Z Notation: A Reference Manual. Prentice Hall, 1989.
- [13] SULLIVAN, K. J., AND NOTKIN, D. Reconciling environment integration and software evolution. ACM Transactions on Software Engineering and Methodology 1, 3 (July 1992), 229-268.

A Z Notation Used in this Paper

The Z notation is a mathematical language developed mainly at the Programming Research Group at the University of Oxford over the last 15 years. The mathematical roots of Z are in first order logic and set theory. The notation uses standard logical connectives (\land , \lor , \Rightarrow , etc.) and settheoretic operations (\in , \cup , \cap , etc.) with their standard semantics. Using the language of Z, we can provide a model of a mathematical object. That these objects bear a resemblance to computational objects reflects the intention that Z be used as a specification language for software engineering. In this appendix, we describe the basics of the Z notation used in this paper. The standard reference for practitioners of Z, and the basis for our use of Z, is Spivey's reference manual [12].

A Z specification consists of sections of mathematical text interspersed with prose. The mathematical text is a collection of types together with some predicates that must hold on the values of each type. Types in Z are sets of values. Z provides some fundamental types in its basic toolkit that are primitive, such as N for natural numbers and Z for integers. In addition, we can introduce further primitive types, called given types, by writing them in square brackets. By convention, given types are written in all capital letters. The construction of elements in a given type is not provided in a specification, usually because that level of detail is not necessary for the purposes of the specification. Prose surrounding the declaration of a given type should indicate the reason the specifier has introduced the type rather than use an existing type. For example, we could introduce two given sets to represent all possible authors and papers that those authors might write. For use in this appendix, no further information about authors or papers need me made explicit, so we write:

[AUTHOR, PAPER]

An element of a type is declared using a colon (:). So we would write *author* : AUTHOR and read this as "*author* is of type AUTHOR", meaning *author* is an element in the set of values defined by AUTHOR. Since AUTHOR is a set, we could also write *author* $\in AUTHOR$, using the set membership function \in . Z uses the : notation when a variable is declared and \in to express predicates over bound variables.

New types can also be defined by constructing them from primitive types using the following type constructors:

- P X is the set of all subsets with elements from type X, also called the powerset of X.
- $X \times Y$ is the type consisting of all ordered pairs (x, y) whose first element is of type X and whose second element is of type Y, also called the cross-product of X and Y.
- seq X is the set of all sequences, or lists, of elements from X, including empty and infinite sequences.
- bag X is the set of all bags of elements from X.
- Relations and functions between types identify special subsets of the cross product type. The ones used in this paper are:
 - $X \leftrightarrow Y$ is the set of all relations between domain type X and range type Y. A relation is simply a subset of $X \times Y$.
 - $X \rightarrow Y$ is the set of all partial functions between X and Y. A partial function does not have to be defined on all elements of its domain type.
 - $X \longrightarrow Y$ is the set of all total functions. Total functions are defined on all elements of the domain type.
 - $X \rightarrow Y$ is the set of all partial functions from X to Y whose inverse is a partial function from Y to X (also called 1-1 or injective).
 - $X \rightarrow Y$ denotes the total injective functions from X to Y.
 - X → Y denotes the bijective functions from X to Y, i.e., the functions from X to Y that are a 1-1 correspondence (total, injective and surjective).

Z has a special type constructor, called the *schema*, an abstract version of the Pascal record or the C struct type constructors. A schema defines a binding of identifiers (or variables) to their values in some type. For example, we could specify the type *Proceedings* as a schema for a typical conference proceedings. The information we might want to specify about a proceedings would be the set of all authors and an index from authors to the papers they wrote. We represent this binding in the boxed schema notation below.

```
Proceedings______
authors : P AUTHOR
```

```
index: AUTHOR \leftrightarrow PAPER
```

A "dot" notation is used to select elements of a schema type. So we could refer to the authors in the proceedings sigsoft93 : Proceedings by writing sigsoft93.authors.

In addition to declaring the bindings between identifiers and values, a schema can specify invariants that must hold between the values of identifiers. In the boxed notation, these invariants are written under a dividing line. All common identifiers below the line are scoped by the declarations above the line. If we wanted to model the invariant that the set of authors in type *Proceedings* can and must include only those authors appearing in the index, we could state that authors is the domain of the *index* relation. We would write this as follows.

EssentialProceedings	
authors : P AUTHOR	
$index : AUTHOR \leftrightarrow PAPER$	
authors = dom index	

Z allows for schema inclusion to facilitate a more modular approach to a specification. In the above example, we could have introduced the invariant on the set of authors as

Essentia Proceedi	lProceedings ngs	 ·····
authors	= dom index	

including the declarations and invariants of *Proceedings* in the new schema *EssentialProceedings*. Z defines a calculus of schema operations of which inclusion is just one example. We do not use many schema operations in this paper, so we direct the interested reader to Spivey's reference manual.

In addition to the schema calculus for defining schema expressions, Z usage relies on some notational conventions for describing the behavior of state machines. The schema represents a binding from identifiers to values. We can view this binding as the static description of some state machine, that is, the view of the state machine at some point in time. Operations on the state machine are transitions from one legal state to another and can be described as a relationship between the values of identifiers before and after the operation. One of the most common conventions is the Δ convention for describing operations. If *Schema* is a schema type, then Δ *Schema* is notationally equivalent to two "copies" of *Schema*, one of which has all of its identifiers decorated with dashes (') to indicate the state after the operation. So, we could write

_ ProceedingsOp	_
$\Delta Proceedings$	

which is equivalent to

ProceedingsOp	
Proceedings	
Proceedings'	
	••••••••••••••••••••••••••••••••••••

or

 $ProceedingsOp _____ authors : \mathbf{P} AUTHOR \\ index : AUTHOR \leftrightarrow PAPER \\ authors' : \mathbf{P} AUTHOR \\ index' : AUTHOR \\ index' : AUTHOR \leftrightarrow PAPER \\ \end{cases}$

Some other operations and notational conventions used in Z are:

- Point == N × N introduces the type Point as a type synonym for the cross product. Type synonyms are a notational convenience.
- If f is a relation, function or sequence, then dom f is the domain of f and ran f is the range of f.
- If S is a set (or sequence), then # S is the size (or length) of S.
- $a \cap b$ is the concatenation of sequences a and b.
- If R is a relation, then R^{\sim} is its relational inverse and R^+ is its transitive closure. If S is a set of elements in the domain type of R, then R(S) is the image over R of the set of elements in S, that is, the set of elements in the range type of R that are related to elements in S under R.
- ∀ decl | pred1 pred2 is read "for all variables in decl satisfying pred1, we have that pred2 holds."
- $\exists decl \mid pred_1 \bullet pred_2$ is read "there exist(s) variable(s) in decl satisfying $pred_1$ such that $pred_2$ holds."
- { decl | pred expression} is a set comprehension for the set of values expression ranging over variables in decl satisfying the predicate pred.