

# Reconciling the Needs of Architectural Description with Object-Modeling Notations

**David Garlan**

Carnegie Mellon University  
Pittsburgh, PA 15213 USA  
garlan@cs.cmu.edu

**Andrew J. Kompanek**

Carnegie Mellon University  
Pittsburgh, PA 15213 USA  
kompanek@cs.cmu.edu

**Pedro Pinto**

Blue Capital Group  
Chapel Hill, NC 27514 USA  
ppinto@cs.cmu.edu

## ABSTRACT

Complex software systems require expressive notations for representing their software architectures. Two competing paths have emerged. One is to use a specialized notation for architecture – or architecture description language (ADL). The other is to adapt a general-purpose modeling notation, such as UML. The latter has a number of benefits including familiarity to developers, close mapping to implementations, and commercial tool support. However, it remains an open question as to how best to use object-oriented notations for architectural description, and, indeed, whether they are sufficiently expressive, as currently defined. In this paper we take a systematic look at these questions, examining the space of possible mappings from ADLs into object notations. Specifically, we describe (a) the principle strategies for representing architectural structure in UML; (b) the benefits and limitations of each strategy; and (c) aspects of architectural description that are intrinsically difficult to model in UML using the strategies.

## Keywords

Software Architecture, Object Modeling, UML

## 1 INTRODUCTION

A critical level of abstraction in the description of a complex system is its software architecture. At an architectural level one describes the principal system components and their pathways of interaction. Architectural descriptions are typically used to provide an intellectually tractable, birds-eye view of a system, and to permit design-time reasoning about system-level concerns, such as performance, reliability, portability, and conformance to external standards and architectural styles.

In practice most architectural descriptions are informal documents. They are usually centered on box-and-line diagrams, with explanatory prose. Visual conventions are idiosyncratic, and usually project-specific. As a result, architectural descriptions are only vaguely understood by developers, they cannot be analyzed for consistency or completeness, they are only hypothetically related to implementations, their properties cannot be enforced as a

system evolves, and they cannot be supported by tools to help software architects with their tasks.

To improve the situation a number of people have suggested the use of more standardized and formal notations for architectural description. Viewed broadly, there are two main sources of such recommendations. One is from the software architecture research community, which has proposed a number of “architecture description languages” (ADLs) specifically designed to represent software and system architectures [16]. These languages have matured considerably over the past five years. For example, most now come with tool sets that support many aspects of architectural design and analysis, such as graphical editing, code generation, run-time monitoring, anomaly detection, and performance analysis.

The other is from the object-oriented community. Object modeling notations have had considerable success both as domain-modeling and implementation-modeling notations. For the former, object-oriented entity-relation diagrams (and related notations) have provided a natural mechanism to represent domain entities and their relations. In the latter case, for systems that are to be implemented in an object-oriented fashion, object-modeling notations provide a natural way to represent class structures, relationships between those classes, and certain aspects of class/system behavior.

But are object-oriented notations also appropriate as architectural descriptions? A number of authors have argued yes, claiming that the current capabilities are precisely what is needed [3, 12]. Moreover, recently there have been a number of proposals that attempt to show how concepts found in ADLs can be mapped directly into a notation like UML [10,15,17].

Unfortunately, each proposal is different, and each covers only *some* aspects of architectural description. We believe that what is needed is more systematic view of the problem, one that clarifies what needs to be modeled at an architectural level, identifies candidate constructs in an object modeling language, and examines the strengths and weaknesses of adopting a particular modeling strategy. Practitioners could then decide more rationally which, if any, of the possible techniques is most appropriate.

In this paper we do just that for the structural aspects of

Draft of November 12, 1999.  
Submitted for Publication.

software architecture. We start by examining the needs for architectural description—based on what is now done informally in practice and what features are captured by ADLs. Then we consider a number of strategies for representing architectures in UML, contrasting the advantages and drawbacks of each. Finally, we summarize the key lessons from this activity.

## 2 RELATED WORK

A number of people have examined ways to model architectures using object notations. Among the earliest of these were purveyors of object-oriented methods, who attempted to provide a uniform path from requirements to implementations using only object notations. In these treatments architecture is usually not specifically called out as a representation in need of special treatment, but simply viewed as a form of high-level object-oriented design.

The lack of explicit treatment of architectures in those methods prompted several authors to take a more detailed look at the needs of architectural modeling. Kruchten [12] for example, proposes a set of (4+1) views to capture the various aspects of architectural information. This approach has been incorporated in the Rational Unified Process.

Other authors have examined the use of object-oriented notations for modeling architectural styles and patterns [3]. These authors have largely confined their attention to the specialized part of the problem of capturing architectural idioms.

More recently there have been a number of attempts to understand more broadly how to *map* architectural models as expressed in an architecture description language into object notations [10,15,17,18]. As we noted earlier, each of these proposes a particular form of modeling that exploits the constraints of the particular architectural domain or language. Unlike this paper, they have not attempted to consider systematically of the *space* of possible embeddings, or enumerate the situations under which *alternative* embeddings might be desirable – although they do serve as relevant data points.

## 3 ARCHITECTURAL DESCRIPTION

The software architecture of a system defines its high-level structure, exposing its gross organization as a collection of interacting components. Currently there is considerable diversity in the ways that practitioners represent architectures, although most depend on informal box-and-line diagrams.

In an effort to put architectural description a more solid notational and semantic footing, over the past decade a number of architecture description languages have been proposed [1,2,6,8,13,14,17,19,27]. While these languages (and their associated tool-sets) differ in many details, there has emerged a general consensus about the main ingredients of architectural description. Focusing on architectural structure, we take that core set of concepts as

the starting point for this paper<sup>1</sup>.

In this shared ontology there are five basic concepts: components, connectors, systems, properties, and styles.

*Components* represent the primary computational elements and data stores of a system. Intuitively, they correspond to the boxes in box-and-line descriptions of software architectures. Typical examples of components include clients, servers, filters, objects, blackboards, and databases. Components may have multiple interfaces (which we will call *ports*), each interface defining a point of interaction between a component and its environment.

*Connectors* represent interactions among components. Computationally speaking, connectors mediate the communication and coordination activities among components. That is, they provide the “glue” for architectural designs, and correspond to the lines in box-and-line descriptions. Examples include simple forms of interaction, such as pipes, procedure call, and event broadcast. Connectors may also represent complex interactions, such as a client-server protocol or a SQL link between a database and an application. Connectors can have interfaces that define the *roles* played by the various participants in the interaction represented by the connector.

*Systems* represent configurations (graphs) of components and connectors. Systems are described by indicating which component ports fill which connector role. In general, systems may be hierarchical: components and connectors may represent subsystems that have *internal* architectures. We will refer to these as *representations*. When a system or part of a system has a representation, it is also necessary to explain the mapping between the internal and external interfaces.

*Properties* represent additional information (beyond structure) about the parts of an architectural description. Although the properties that can be expressed in different ADLs vary considerably, typically they are used to represent anticipated or required extra-functional aspects of an architectural design. For example, some ADLs allow one to calculate overall system throughput and latency based on performance estimates of the constituent components and connectors [28]. In general, it is desirable to be able to associate properties with any architectural element in a description (components, connectors, systems, and their interfaces). For example, an interface (port or role) may describe an interaction protocol.

*Types and Styles* represent families of related systems. An

---

<sup>1</sup> These concepts are essentially those in Acme [7], a community-defined architecture interchange language.

architectural *style* typically defines a vocabulary of design element types as a set of component, connector, port and role types, together with rules for composing instances of the types [9]. Examples include data-flow architectures based on graphs of pipes and filters, blackboard architectures, and layered systems.<sup>2</sup>

To illustrate the use of these concepts, consider the simple example in Figure 1. The system is described in a pipe-filter style, which provides a design vocabulary consisting of a filter component type and pipe connector type, input and output port types, and a port-port binding type. In addition there would likely be constraints (not shown here) that ensure, for example, that the reader/writer roles of the pipe are associated with appropriate (input/output) ports. The system itself defines a simple string-processing application that extracts and sorts text. The system is described hierarchically: the *MergeAndSort* component is defined by a representation that is itself a pipe-filter system. Properties of the components and connectors (not illustrated) might include performance characteristics used by a tool to calculate system throughput.

#### 4 UML BACKGROUND

UML unifies a number of object modeling notations in a common framework and is becoming the standard object notation for industry [4,5,24,25]. While a detailed description of UML is beyond the scope of this paper, we summarize its principal constructs (known as model *elements*) that can be used to model software architectures:

*Classes, Interfaces and Objects* Classes are the primary construct used for describing the logical view of a system. They are used to capture both domain-level abstractions as well as the abstractions in an object-oriented software design. Classes have properties in the form of *attributes*, publish abstract services in the form of *operations*, and are logically related to one another using *associations*. Classes may expose their functionality through a set of supported *interfaces*, collections of related operations. Classes have instances in the form of *objects*, which are often used in models that depict a behavior under a particular scenario.

*Component and Component Instances* UML can also be used to describe the physical structure of a software system in terms of the layout of its physical parts, called *components*, computing resources, called *nodes* and bindings of components to resources. Like classes, components in UML expose their functionality through




Figure 1: A system in the Pipes and Filters style

interfaces. Components are typically related to each other using *dependency* relationships.

*Packages and Package Instances* UML provides grouping and composition mechanisms that are used to partition large UML models into manageable chunks called *packages*. UML defines a type of package called a *sub-system*, which is typically used to encapsulate the object models that define a system.

*Relationships* Model elements are related to one another with *associations* and *dependencies*. Dependency is the most generic relationship in UML, indicating that an element depends in some way on the definition of another element. Association is a richer relationship that describes an abstract relationship between classes and the roles the classes play in the relationship.

*Stereotypes* To allow the extension of UML with domain-specific concepts, UML provides a mechanism for associating constraints with elements of a model, using a constraint language, *OCL* [23]. These constraints can be grouped and named using a construct called a *stereotype*. UML also includes a set of standard stereotypes. For example, the `<<process>>` stereotype for classes and the `<<refines>>` stereotype for dependencies.

The above constructs can be composed in various ways in a UML model and visualized in diagrams. The constructs available for description varies among these elements, although arbitrary textual *annotations* may be associated with any of them. Frequently, these annotations are in the form of *tagged values* – or arbitrary attribute-value pairs.

UML also defines a set of models for describing the dynamic behavior of a system, including collaboration diagrams that specify system behavior using event-based *interaction scenarios*, descriptions based on *state machines*, and *use cases*.

#### 5 STRATEGIES & EVALUATION CRITERIA

We now consider five strategies for modeling architectural structure (as expressible in ADLs) using UML. We will organize the presentations of alternatives around the choices for representing component types and component instances, since the components are typically the central design elements of an architectural description. For each choice we then consider sub-alternatives for the other architectural elements (connectors, styles, etc.). Of the five strategies, the first three consider ways to use classes and objects to model components. The fourth considers an approach based on UML components. The fifth considers the suitability of UML packages.

<sup>2</sup> Various terms are used to distinguish different kinds of families of architectures: “product-line frameworks”, “component integration standards”, “architectural patterns”, “styles”, and “idioms”.

However, before considering those strategies, we need to be clear about the criteria for evaluation. Ideally we would like a mapping strategy to exhibit three characteristics:

1. **Semantic Match:** It should respect documented UML semantics and the intuitions of UML modelers. The interpretation of UML model should be a close to the interpretation of the original ADL description so the model is intelligible to both designers and UML-based tools. In addition the mapping should produce legal UML models.
2. **Legibility:** The resulting architectural descriptions in UML should bring conceptual clarity to a system design, avoid unnecessary clutter and highlight key design details.
3. **Completeness:** The architectural concepts identified in Section 3 should be representable in UML model.

## 6 CLASSES & OBJECTS: VARIATION #1, TYPES AS CLASSES AND INSTANCES AS OBJECTS

Perhaps the most natural candidate for representing component and connector types in UML is the class concept. Class diagrams in UML are the main way in which the logical view of a system is described. Classes describe the conceptual vocabulary of a system just as component and connector types form the conceptual vocabulary of an architectural description in a particular style. Additionally, the relation between classes and instances has an appealing similarity to the relationship between architectural types and their instances. The mapping described for the C2 architectural style in [15] is a variation on this approach.

Figure 2 illustrates the general idea. Here we characterize the filter architectural type as the UML class “filter”. Instances of filters, such as “split” are represented as corresponding objects in an object (instance) diagram.

We now take a closer look at this strategy by examining how the concepts in Section 3 can be described in UML.

### Components

As noted, the type/instance relationship in architectural descriptions is a close match to the class/object relationship in a UML model. In UML, classes, like component types in architectural descriptions, are first-class entities and are the richest structures available for capturing software




Figure 2 : Types as classes, instances as objects (simplest variation).

abstractions. The full set of UML descriptive mechanisms are available to describe the structure, properties and behavior of a class, making this a good choice for depicting detail and doing analysis using UML-based analysis tools.

Properties of architectural components can be represented as class attributes or with associations; behavior can be described using UML behavioral models; the <<instanceOf>> dependency can be used to relate a component instance to its type(s) and generalization can be used to relate a set of component types. The semantics of an instance or type can also be elaborated by attaching one of the standard stereotypes (e.g., indicating that a component runs as a separate process with the <<process>> stereotype).

However, the match between component types and classes, and component instances and objects, is not perfect. In architectural descriptions a component instance often has additional structure not required by the component’s type. For example, a component instance might define additional ports not required by its type, or associate an implementation in the form of additional structure that is not part of its type’s definition. However, in UML an object cannot include parts that its class does not also define.

### Ports

There are six reasonable ways to represent ports.

#### Option 1: No explicit representation

We can leave ports out of the UML model entirely (as illustrated above). This choice leads to the simplest diagrams, but suffers from the obvious problem that there is no way to characterize the names or properties of the ports. However, this might be a reasonable choice under certain restricted situations, such as when components have only a single port, or when the ports can be inferred from the system topology.

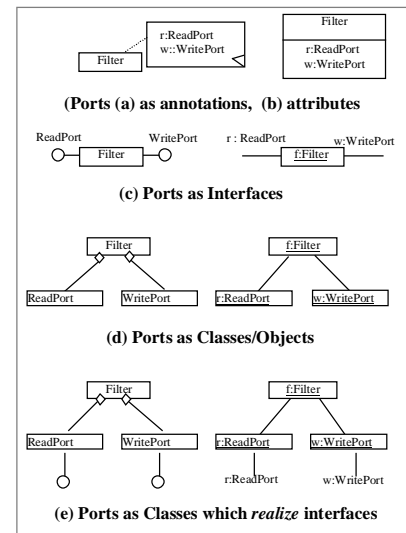


Figure 3: Representing ports

#### Option 2: Ports as Annotations

We can represent ports as annotations (Figure 3a). This approach provides a home for information about ports, although annotations have no semantic value in UML

and hence cannot be used as a basis for analysis. Again, if the detailed properties of a port are not of concern this might be a reasonable approach.

#### *Option 3: Ports as Class/Object Attributes*

Ports can be treated as attributes of a class/object. In this approach ports are part of the formal structural model, but they can have only a very simple representation in a class diagram – essentially a name and type (see Figure 3b).

#### *Option 4: Ports as Interfaces*

Describing port types as UML interfaces has three advantages. First, the interface and port concepts have a similar meaning: they both represent the way in which an entity can interact with its environment. Second, the UML “lollipop” notation provides a compact description of a port in a class diagram depicting a component type. In an instance diagram, a UML association role (corresponding to a port instance) qualified by the interface name (the port type) provides a compact way to designate that a component instance is interacting through a particular port instance. Finally, this approach provides visually distinct depictions of components and ports, in which ports can clearly be seen as subservient to components (see Figure 3c).

However, while the interface and port concepts are superficially similar, they are not identical. An interface passively exposes a set of operations offered by classes that can be invoked by the environment. In contrast, the description of a port in an ADL often includes both the services provided by the component, as well as those it requires from its environment. Furthermore, it is meaningful for a component type to have several instances of the same port type, while it is not meaningful to say that a class realizes several versions of the same interface (let alone indicate that a port may have a variable number of a given port type). For example, there is no way to define a “splitter” filter type that has two output ports of the same “type” using this technique. Finally, unlike classes, interfaces do not have attributes or substructure.

#### *Option 5: Ports as Classes*

Another alternative is to describe ports as classes contained by a component type (see Figure 3d). This overcomes the lack of expressiveness of UML interfaces (option 4). We can now represent port substructure and indicate that a component type has several ports of the same type. A component instance is modeled as an object associated with a set of port objects. Unfortunately, this approach also suffers from problems: by representing ports as classes, we not only clutter the diagram, but also lose clear discrimination between ports and components. It is possible to make the diagram more suggestive using a variation on this scheme in which the ports are contained classes. But then indicating points of interaction is somewhat counterintuitive, as containment usually indicates that a class owns other classes whose instances may or *may not* be accessible through instances of the parent class.

#### *Option 6: Ports as UML Classes that realize interfaces*

The final option is a combination of options 4 and 5 (see Figure 3e). We represent ports as classes which themselves expose interfaces. While this option more expressive than the previous two techniques, it suffers from the semantic mismatch problem of the previous two approaches. It is also the least visually appealing. Unless ports are discriminated from components visually, in a complete system, the added clutter could mask the overall topology of components, defeating one of the main purposes of architectural description. It also makes a diagram more difficult to interpret because a reader is expected to understand that a pair of objects stands for a single object in the original model.

### **Connectors**

#### *Option 1: Types as Associations; Connectors as Links*

In an architectural box-and-line diagram of a system, the lines between components are connectors. One tempting way to represent connectors in UML is as associations between classes or links between objects (see Figure 3). The approach is visually simple, provides a clear distinction between components and connectors, and makes use of the most familiar relationship in UML class diagrams: association. Moreover, associations can be labeled, and when a direction is associated with the connector it can be indicated with an arrow in UML.

Unfortunately, although the identification between connectors and associations is visually appealing, connectors have a different meaning than associations. A system in an architectural description is built-up by choosing components with behavior exposed through their ports and connecting them with connectors that coordinate their behaviors. A system’s behavior is defined as the collective behavior of a set of components whose interaction is defined and limited by the connections between them. In contrast, while an association or link in UML represents a potential for interaction between the elements it relates, the association mechanism is primarily a way of describing a conceptual relationship between two concepts.

Using an association to represent architectural connection has other problems. Since associations are relationships *between* UML elements, an association cannot stand on its own in a UML model. Consequently, there is no way to represent a connector type in isolation. To have any notion of a connector type within this scheme, one must resort to naming conventions or the use of stereotypes whose meaning is captured by an OCL description to indicate connector type. Also, the approach does not allow one to specify the interfaces to the connector (i.e., the roles).

If links are chosen to represent connector instances, the UML elements that are associated in a model depend on the choice of port representation. If ports are not represented in the model, a connector will be a link between comp-

onents. If ports are represented as objects in an instance diagram, the link will be between the port objects. If ports are described with interfaces there is a problem – since interfaces cannot be related by association, only by dependency.

#### *Option 2: Types as Association Classes*

One solution to the lack of expressiveness of an association is to qualify the association with a class that represents the connector type. This way the attributes of a connector type or connector can be captured as attributes of a class or object. Unfortunately, this technique still does not provide any way of explicitly representing connector roles.

#### *Option 3: Types as Classes; Connectors as Objects*

To provide an explicit representation of roles, we must represent connector types as classes just as we represented component types. We have the same options for representing roles as we had for ports: roles as interfaces realized by a class, as “child” classes contained by a class or as child classes that realize interfaces. The different approaches for representing roles have the same consequences as the analogous choices for representing ports: increasing expressiveness at the expense of visual simplicity.

Given a choice for representing ports, and a choice for representing roles, there are two options for representing an attachment. 1. As association/link between elements that represent ports/roles (for classes, objects). 2. As dependency between elements that represent ports/roles.

### **Systems**

#### *Option 1: Systems as UML Packages*

The primary mechanism in UML for grouping related elements is the package. In fact, UML defines a standard package stereotype called <<subsystem>> to group UML models that represent a logical part of a system. The choice of packages to model systems is appropriate for *any* choice of mappings of components and connectors and works particularly well for grouping classes. Packages also have a visual appeal since elements can be depicted as contained in a package.

Unfortunately, subsystems have different semantics than systems in an architectural description. In a model, a package represents set of elements that may be imported into another context, but not a structure per se. In contrast, a system in architectural design is a structure with sub-parts in the form of its components and connectors. Unlike classes, packages also lack attributes for defining system level properties.

#### *Option 2: Systems as Objects*

A second option is to use objects to represent systems. If architectural instances are represented as objects an explicit system class can be introduced whose instances contain the component and connector objects that make up the system. Then, we can capture richer semantics using attributes and associations/links.

Unfortunately, this approach has the significant drawback that by representing a system in the same as components and connectors we lose the semantic distinction between a system as a configuration of elements. It is also introduces visual clutter.

#### *Option 3: Systems as Packages containing a System object*

The final alternative is to combine the first two techniques, by representing a system as a package that contains a class that describes the system. This combines expressiveness of option 2 with the visual advantages of option 1. However, we still have the basic semantic mismatch and the additional clutter complicates a diagram and makes interpretation of the diagram more difficult.

There are of course many variations on the last two options, corresponding to the same variations on the use of classes, objects and stereotypes to describe components which we described in the last section and in the next.

### **Styles**

We can use the same mechanisms for describing styles. If we represent a style as a package we can stereotype it <<framework>> (a standard UML stereotype), and also take advantage of the support for defining design patterns using templates. We can also capture the relationship between styles and instances using dependency relationships. Similarly, if systems are described with objects, styles can be represented (awkwardly) as classes.

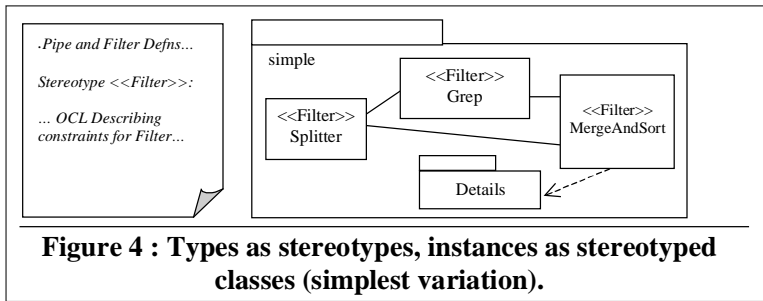
### **Representations**

A representation is just a system that is related to a component or connector and includes bindings that describe the relationship between the element and (one of) its representations. If systems are represented as packages, an element can be related to its representations through dependency, possibly with an annotation indicating the name of the representation. There are a number of options for representing bindings: (1) as annotations (a weak depiction); (2) as dependency between an outer element and inner element (only possible with techniques in which ports and roles are explicitly represented); (3) as set of elements with attributes that describe binding (visually cluttered).

Alternatively, if systems are represented as classes, representations can be associated with their parent element using associations and links.

## **7 CLASSES AND OBJECTS: VARIATION #2, COMPONENT TYPES AS STEREOTYPES AND INSTANCES AS CLASSES**

The second major alternative for modeling component types in UML is to define a special stereotype (using OCL) to represent architectural component types. In this way, we can describe the meaning of our architectural vocabulary in a way that distinguishes an architectural element type concept from the UML class concept. A component instance is then represented as a class with a stereotype.



**Figure 4 : Types as stereotypes, instances as stereotyped classes (simplest variation).**

Using this approach, architectural concepts are clearly distinct from the built-in UML concepts, and in principal, a UML-based modeling environment can be extended to support the visualization and analysis of new architectural types within a style and enforce design constraints captured in OCL. This is essentially the approach taken in [17].

In Figure 4, the Filter Type is defined by a set of constraints expressed in OCL that are identified with the <<Filter>> Stereotype. Filter instances (e.g., splitter) are represented as classes that bear the Filter stereotype.

We now examine this approach in more detail.

### Components

Representing instances as classes has a number of advantages. For example, we can capture a set of possible runtime configurations (e.g., one to many relationships) in a single diagram. By using classes, we also allow component instances to include structure in addition to the structure defined by their types, overcoming a limitation of the class/instance approach.

Unfortunately, the approach has a number of disadvantages. Stereotypes are not first class, so we can't define structure, attributes, or behavior except by writing OCL constraints. Furthermore, there is no way to visualize stereotypes in diagrams, unless, in the future, there is support in UML environments for manipulating, visualizing and analyzing extensions to the UML meta-model. There is also no way to express sub-typing relationships between stereotypes. Among other consequences, this means we can't take advantage of potential analysis capabilities offered by UML tools to analyze architectural types, or associate behavioral models with types.

Furthermore, a class may have only one stereotype. Consequently, we can't use any the other built-in stereotypes to add meaning to a component that already has a stereotype. Arguably, using a *class* to represent an architectural *instance* is also not intuitive.

Using this approach, there are a number of options for representing ports. The ports defined by a component type would be represented as OCL expressions that state what structure a class standing for a component of this type should have. We can represent the ports belonging to a component

*instance* in the same ways we modeled a component type's ports in Variation #1.

### Systems, Representations, and Styles

The same options exist for describing overall systems and for indicating representations as for the first variation, although in this we represent system instances as represented styles in the previous approach. In this case, there can be no explicit representation of a style using UML model elements. Instead, the style is embodied in a set of stereotypes (e.g., <<Filter>>).

## 8 CLASSES AND OBJECTS: VARIATION #3, COMPONENT TYPES AS CLASSES AND INSTANCES AS CLASSES

The third class-based approach is to represent component types as classes, as we did in Variation #1, and component instances as classes as in Variation #2. The observations we made in the previous two sections about the suitability of the classes to represent component types and instances apply here as well: by representing both component types and instances as classes, we have the full set of UML features to describe both component types and instances. We can also capture patterns at both the type (as part of a description of an architectural style) and instance level, supporting the description of a dynamic architecture whose structure evolves at run-time.


Figure 4 illustrates this approach, defining both the "filter" type and instances of filters (e.g., "splitter") as classes.

We now examine this approach in light of the previous two class-based approaches.

### Components

Although, the approach has many of the strengths of the previous two approaches, it suffers from a number of problems. Unlike the second approach, we no longer explicitly distinguish architectural vocabulary and the UML class concept. Representing both types and instances as classes also blurs the distinction between type and instance, stereotypes may be used to reinforce this distinction.

There are two viable options for relating a component type to its instances within this scheme: generalization and dependency. Generalization captures the structural relationship between a type and instance (namely, the instance "subclass" must be substitutable for the type) but it blurs



**Figure 5 : Types & instances as classes (simplest variation).**

the type-instance distinction. The other option is to represent this relationship as a (perhaps stereotyped) dependency, which is semantically weaker but carries less semantic baggage than the generalization relationship.

### Connectors

The same issues arise when describing connections. For a full discussion of options for describing connections in terms of classes, see variation #1.

### Systems, Representations and Styles

The options for representing systems and styles are similar to those in the two previous approaches using packages or classes and objects.

## 9 USING UML COMPONENTS

UML includes a “component” structure in the model used to describe the *implementation* artifacts of a system and their deployment. As such, it is an alternative for describing architectural components. A component diagram is often used to depict the topology of a system at a high level of granularity and plays a similar function, although at the implementation level, as an architectural description of a system. In Figure 6 we represent the “filter” type as a UML component and instances (e.g., “splitter”) as instances of this UML component.

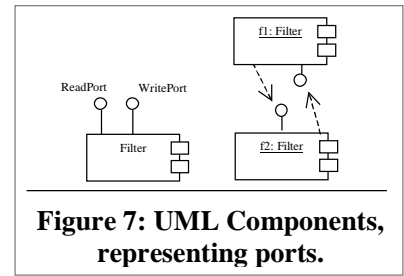
### Components

UML components are a natural candidate for representing architectural components. Components have interfaces, they can be deployed on hardware, and they commonly carry a stereotype and a custom visualization to reflect the sort of component being depicted. UML components are often used as part of diagrams that depict an overall topology, and just as it is natural to talk about mapping architectural components to hardware, components are assigned to nodes in UML deployment diagrams. For some architectural styles, the identification of abstract components with implementation-level components is a reasonable choice.

Unfortunately, in UML components are defined as concrete chunks of implementation (e.g., executables, or dynamic link libraries) that realize abstract interfaces – unlike components (and component types), which frequently have only an indirect relationship to deployable piece of a system. Nonetheless, the concepts share more than a name. Components expose interfaces just as classes do, and can be used to represent the ports exposed by a component, just as they were used in the strategy based on classes and

objects.

However, the rich set of class associations available to relate classes are not available for components, limiting how we can describe sub-structure (i.e., ports), represent patterns as part of a style description and show connection. Moreover, UML behavioral models cannot reference components.



**Figure 7: UML Components, representing ports.**

### Connectors

There are two natural choices for representing connectors in this scheme: as dependencies between a component the ports/interfaces realized by a component (visually simple but lacking expressiveness) or as components themselves. If we represent connector instances as dependencies between components, we have the option of representing connector types as stereotypes, with consequences we addressed in previous sections. Unfortunately, although dependencies are visually appealing, the built-in dependency notion in UML does not adequately capture the notion of architectural connection or provide an explicit descriptive capability. Representing a connector as a UML component blurs the distinction between components and connectors, a problem (partially) addressed by introducing new stereotypes.

### Systems, Representations, and Styles

The only reasonable way to describe systems and styles with UML component models is as packages since this is the only grouping mechanism available for components.


## 10 USING PACKAGES

We now describe how a subsystem (a stereotyped package) can be used to describe the components in a system and the component types in a style. This approach has the appeal that packages are an ideal way to describe coarse-grained elements as a set of UML models. The package construct is already familiar to UML modelers, and to those extending UML, as a way of bundling large pieces or views of a system. The approach used in [10] is based on packages.

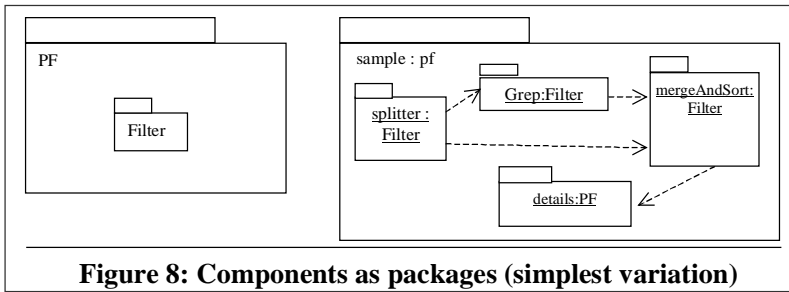
In Figure 8, we describe the “filter” type as a package and filter instances as package instances (e.g., “splitter”).

### Components

The subsystem construct is used in UML to group or encapsulate a set of model elements that describe some logical piece of a system, similar to components in architectural description. Subsystems (indeed, any package) can include structures based on any of the UML models. This has an advantage over describing components and connectors as classes – by identifying a component or connector with a package, we can



**Figure 6 : Components as UML components (simplest variation).**



**Figure 8: Components as packages (simplest variation)**

not only include structure as classes (or objects), we can also include a behavioral models. There are many options for describing a component or connectors substructure. It also has a visual appeal – substructure can be visualized as “embedded” in the package. Component and Component types would be modeled in essentially the same way, although one could also take advantage of the UML template mechanism when defining a type.

Although visually appealing, this approach suffers from a number of problems. In UML, a package is (semantically speaking) little more than a folder in which related models and elements can be stored. One does not talk about a package having interfaces; instead a package makes its elements (e.g., classes) available to other packages through export. Representing substructure (like ports) as elements contained by package is also counterintuitive. The fact that certain elements correspond to ports, others to properties, others to reps, is misleading. Expressing the behavior of a component would also be awkward, since packages themselves do not have dynamic behavior in UML. Packages may only be related by dependence, which restricts how connection/attachment can be depicted. It also blurs the distinction between a system and a component in a system.

### Connectors

There are two natural choices for representing connectors in this scheme: dependencies (visually simple but lacking expressiveness), or packages themselves. While dependencies have visual appeal, the dependency notion does not fully capture the notion architectural connection. As we’ve noted before, a package does not have to describe of connector properties and structure. As with components, we could include behavioral description of the connector within the package, although relating it meaningfully to the package itself is difficult, as it was for components.

### Systems, Representations, and Styles

The only mechanism available for grouping together related packages is the package mechanism itself, which we considered in previous sections. To distinguish packages representing systems and styles from those representing components or connectors, we can define new stereotypes.

## 11 DISCUSSION AND CONCLUSIONS


We have briefly examined five strategies for encoding in UML the representations of architectural structure

expressible using modern ADLs: for each of these we considered a number of variations and evaluated the strategies with respect to completeness, legibility, and semantic match. What are the lessons to be gained from this?

First, *there is no single best way to encode ADLs in UML*. Each of the strategies has strengths and weaknesses, depending on how well they support the evaluation criteria. With respect to completeness and legibility there is a typically a tradeoff: encodings that emphasize completeness (by providing a semantic home for all of the aspects of architectural design) tend to be verbose, while graphically appealing encodings tend to be incomplete. Hence, the best strategy will depend on what aspects of architectural design needed to be represented. In restricted situations (for example, if there is only one type of connector) it may be preferable to use an incomplete, but visually appealing, encoding.

Second, *all of the encodings exhibit some form of semantic mismatch*. UML’s vocabulary of classes, objects, packages, components, associations, etc., while varied and rich, is ultimately designed to support an object-oriented view of software design. That view emphasizes composition of systems from parts that provide certain services, and depend on the services of other components. On the other hand, in its most general form, architecture takes a broader view, particularly with respect to component interaction, allowing coordination and packaging that depend on richer forms of interaction beyond service invocation. Hence, it is likely that none of the encodings will be completely satisfying for many architects.

There are other problems with the use of UML for representing architecture. Perhaps the most serious is that by shifting from special-purpose notations for architecture to a general-purpose design notation, one gives up several capabilities. First, is analyzability. Virtually all ADLs support some form of non-trivial analysis (or code generation.. Analysis is possible because of the constraints embodied in the ADL (for example, requiring certain kinds of semantic information that can be processed by tools).



**Figure 9: Refinement Paths**

Second, is guidance. By limiting the design vocabulary to architectural concepts, ADLs encourage an architectural approach to design. In a more general language like UML it is easy to mix higher-level architectural issues with lower-level ones. (This is analogous to using C as an

object-oriented programming language.)

Given these observations, one might well ask whether there

are reasonable alternatives to direct encoding of architecture in ADL. We see two plausible paths:

1. Continue to use ADLs, but map to OO implementations. In this approach architecture description retains its own notations, but tools are provided to convert those descriptions to *lower-level* object notations in situations where the implementation is done in an object-oriented fashion. Referring to Figure 9, we follow path A- D, rather than A-C-E (or A-B-E, if we are staying completely within the OO universe).
2. Extend UML to include an explicit architectural vocabulary. This could be done by extending the UML meta-model, or by defining a *Profile* for architectural design<sup>3</sup>. Indeed, we can already see the inclusion of architectural notions, such as ports, in proposals for real-time extensions to UML [21,26].

Both of these are promising avenues of future work. Additionally, we see considerable value in extending our examination of mappings to non-structural aspects of software architecture, such as behavior, performance, and reliability. As with structure, several authors have examined possible encodings, although not in a comparative study. Finally, to make more progress in reconciling architecture description with UML it will be important to consider the problem at a more formal level. In this paper we were forced to appeal to intuition regarding issues of semantic similarity. It would be nice to have a more precise foundation on which to base this comparison.

## REFERENCES

1. Allen, R.; Garlan, A. Formalizing Architectural Connection. *Proceedings of the 16<sup>th</sup> Intl. Conf. on SW Eng.*, 1994.
2. Binns, P.; and Vestal, S. Formal Real-Time Architecture Specification and Analysis. *Proceedings of 10<sup>th</sup> IEEE Wkshp on Real-Time OS and Sw*, May 1993.
3. Buschmann, F; Meunier, R.; Rohnert, H.; Sommerlad, P. and Stal, M. *Pattern-Oriented Software Architecture, A System of Patterns*. John Wiley & Sons, 1996.
4. Booch, G.; Rumbaugh, J and Jacobson, I. *The Unified Modeling Language User Guide*. Addison-Wesley.
5. Booch, G.; Rumbaugh, J and Jacobson, I. *The Unified Modeling Language Reference Manual*. Addison-Wesley.
6. Coglianese, L. and R. Szymanski. DSSA-ADAGE: An Environment for Architecture-based Avionics Development. *In Proceedings of AGARD'93*, May 1993.
7. Garlan, D.; Monroe, R.T. and Wile, D. *Acme Reference Guide*. Available from <http://www.cs.cmu.edu/~acme>.
8. Garlan, D.; Allen, R.; and Ockerbloom, J. (1994), Exploiting Style in Architectural Design Environments. *SIGSOFT'94*.
9. Garlan, D. and Shaw, M. *Software Architecture: Perspectives on an Emerging Discipline*. Prentice Hall, 1996.
10. Hofmeister, C., R.L. Nord, and D. Soni. Describing Software Architecture with UML. *Proceedings of the TC2 1<sup>st</sup> Working IFIP Conf. on Sw Architecture (WICSA1)* 1999.
11. Kobryn, C. Modeling Enterprise Software Architectures Using UML. *In 1998 Proceedings International Enterprise Distributed Object Computing Workshop*, IEEE, 1998.
12. Kruchten, P.B. The 4+1 View Model of Architecture. *IEEE Software*, pp. 42-50, November 1995.
13. Luckham, D.; Augustin, L.M.; Kenney, J.J.; Vera, J.; Bryan, D.; Mann, W. (1995) Specification and Analysis of System architecture using Rapide. *IEEE Trans on Software Eng.*
14. Magee, J.; Dulay, N.; Eisenbach, S. and Kramer, J. *Specifying Distributed Software Architectures. Proceedings of the 5<sup>th</sup> European Software Eng. Conf.*, 1995.
15. Medvidovic, N. and Rosenblum, S. Assessing the Suitability of a Standard Design Method for Modeling Software Architectures. *Proceedings of the TC2 First Working IFIP Conference on Software Architecture (WICSA1)*, 1999.
16. Medvidovic, N. and Taylor, R.N (1997), A Framework for Classifying and Comparing Architecture Description Languages. *In Proceedings of the 6<sup>th</sup> European Software Engineering Conference together with FSE4*.
17. Medvidovic, N.; Oreizy, P.; Robbins, J.E. and Taylor, R.N. (1996), Using object-oriented typing to support architectural design in the C2 style. *Proceedings of ACM SIGSOFT'96: 4th Symp. on the Found. of Software Eng. (FSE4)*.
18. Monroe, R.T.; Kompanek, A; Melton, R. and Garlan, D. Architectural Style, Design Patterns, and Objects. *IEEE Software*, January 1997.
19. Moriconi, M.; Qian, X.; and Riemenschneider, R. (1995), Correct Architecture Refinement. *IEEE Transactions on Software Engineering*, pp. 356-372 (April).
20. Object Management Group (OMG), Analysis and Design Platform Task Force. White Paper on the Profile mechanism Version 1.0. OMG Document ad/99-04-97, April 1999.
21. OMG. RFP-9, UML Profile for Performance, Scheduling and Time. OMG Document ad/99-93-13.
22. Robbins, R.E.; Medvidovic, D.F.; Redmiles, D.F. and Rosenblum, D.S. Integrating Architecture Description Languages with a Standard Design Method. *In Proceedings of the 20<sup>th</sup> Intl. Conf. on Software Eng. (ICSE'98)*.
23. Rational Software Corporation and IBM (1997), OCL specification. OMG document ad/97-8-08. Available from <http://www.omg.org/docs/ad>.
24. UML Semantics. OMG document ad/97-08-04 (September). Available from <http://www.omg.org/docs/ad/97-08-04.pdf>.
25. UML Notation Guide. OMG document ad/97-08-05. <http://www.omg.org/docs/ad/97-08-05.pdf>
26. Selic, B; Rumbaugh, J. Using UML for Modeling Complex Real-Time Systems. White Paper.
27. Shaw, M.; DeLine, R.; Klein, D.; Ross, T; Young, D.; and Zelesnik, G. Abstractions for Software Architecture and Tools to Support Them. *IEEE Trans. on Software Eng*, 1995.
28. B. Spitznagel and D. Garlan. Architecture-Based Performance Analysis. *Proceedings of the 10th Intl. Conf. on Software Eng. and Knowledge Eng. (SEKE'98)*, 1998.

<sup>3</sup> Profiles are community-based, domain-specific, standardized UML extensions, blessed by the OMG [20].

