Published in *1996 International Conference on Software Reuse*, Orlando, Florida

Validating Component Compositions in Software System Generators

Don Batory and Bart J. Geraci
Department of Computer Sciences
The University of Texas, Austin, Texas 78712
{batory, geraci}@cs.utexas.edu

Abstract[1]

Generators synthesize software systems by composing components from reuse libraries. In general, not all syntactically correct compositions are semantically correct. In this paper, we present domain-independent algorithms for the GenVoca model of software generators to validate component compositions. Our work relies on attribute grammars and offers powerful debugging capabilities with explanation-based error reporting. We illustrate our approach by showing how compositions are debugged by a GenVoca generator for container data structures.

1 Introduction

Software system generators automate the development of software for large families of applications. Generators automatically transform compact, high-level specifications of target systems into actual source code, and rely on libraries of parameterized, plug-compatible, and reusable components for code synthesis.

Generators [Bla91, Bat92a, Bax92, Gom94, Lei94, Nin94] are among many approaches that are being explored to construct customized software systems quickly and inexpensively from reuse libraries. CORBA and its variants simplify the task of building distributed applications from components [Ude94]; CORBA can integrate components that are independently designed and stand-alone modules or executables in a heterogeneous environment. In contrast, generators are closer to toolkits [Gri94], object-oriented frameworks [Joh92], and other reuse-driven approaches (e.g, [Wei90, Sit94]), because they focus on software domains whose components are not stand-alone, that are designed to be plug-compatible and interoperable with other components, and that are written in a single language. The particular class of generators that we consider in this paper, called *GenVoca generators* [Bat92a], is distinguished from the above approaches in that their components are parameterized program transformations that encapsulate consistent data and operation refinements. Components also encapsulate logic to automate domain-specific decisions about when to use a particular algorithm or when to apply a domain-specific optimization. For many domains, such decisions are essential for generating efficient code.

A fundamental problem for all component-based software development technologies is: does a composition of components meet the behavioral (or functional) specifications of the target system? For the case of GenVoca generators, this is the problem of *design rule checking*, i.e., the detection of illegal combinations of components. To be viable tools of future software development environments, it is critical that generators validate component compositions automatically (and suggest repairs when errors are detected), rather than burdening users with the impossible task of debugging generated code.

In this paper, we present domain-independent algorithms for solving the problem of design rule checking for GenVoca generators, and present the domain-specific variants that we have used in the Genesis and P2 projects. Our work is related to Perry's Inscape environment, which (among other topics) dealt with consistency checking in software composition models [Per87-89b]. We adapt and generalize the component consistency checking approach of Inscape to exploit the semantics of layers in the construction of hierarchical software systems. We explain how GenVoca models of software domains are grammars, where sentences correspond to component compositions. By encoding component properties as inherited and synthesized attributes, we find that attribute grammars provide a natural formulation of the legal sentences (component compositions, software systems) of a domain. We illustrate our results by explaining how the P2 data structure generator validates component compositions.

1. This work was supported in part by Microsoft, Schlumberger, the University of Texas Applied Research Labs, and the U.S. Department of Defense Advanced Research Projects Agency in cooperation with the U.S. Wright Laboratory Avionics Directorate under contract F33615-91C-1788.

2 The GenVoca Model

GenVoca is a domain-independent model for defining scalable families of hierarchical systems from components. Its basic premise is that standardizing both the fundamental abstractions of mature software domains *and* their implementations, one can define plug-compatible and interchangeable software "building blocks". Although the number of fundamental abstractions in a domain is rather small, there is a huge number of potential implementations. GenVoca also advocates a layered decomposition of implementations, where each layer or component encapsulates a primitive domain feature. The advantage of GenVoca is *scalability* [Bat93, Big94]: component libraries are relatively small and grow at the rate new components are entered, whereas the number of possible *combinations* of components (i.e., distinct software systems in the domain that can be defined) grows geometrically. Generators that use GenVoca organizations have been built for the domains of avionics, data structures, databases, file systems, and network protocols [Cog93, Bat93, Hei93, Hut91].

Components and Realms. A hierarchical software system is defined by a series of progressively more abstract virtual machines. A *component* or *layer* is an implementation of a virtual machine. The set of all components that implement the same virtual machine is called a *realm*; effectively, a realm is a library of plug-compatible and interchangeable components. In Figure 1a, realms `S` and `T` have three components, whereas realm `W` has four.

```
(a)   S = { a, b, c }

      T = { d[S], e[S], f[S] }

      W = { n[W], m[W], p, q[T,S] }
---------------------------------------
(b)   S :=  a | b | c ;

      T :=  d S | e S | f S ;

      W :=  n W | m W |  p  | q T S ;
```

Figure 1. Realms, Components, and Grammars

Parameters and Transformations. A component has a (realm) parameter for every realm interface that it imports. All components of realm `T`, for example, have a single parameter of realm `S`.[2] This means that every component of `T` exports the virtual machine interface of `T` and imports the virtual machine interface of `S`. Thus, each `T` component encapsulates a mapping or *transformation* between the virtual machines `T` and `S`. Such transformations often involve domain-specific optimizations and the automated selection of appropriate algorithms.

Systems and Type Equations. A software *system* is modeled by a composition of components called a *type equation*. Consider the following two equations:

```
System_1 = d[ b ];
System_2 = f[ a ];
```

`System_1` is a composition of component `d` with `b`; `System_2` composes `f` with `a`. Note that both systems are equations of type `T` (because the outermost component of both systems are of type `T`). This means that both implement the same virtual machine and hence, `System_1` and `System_2` are interchangeable implementations of the interface of `T` (with respect to functionality, not performance).[3]

Grammars, Families of Systems, and Scalability. Realms and their components define a grammar whose sentences are software systems. Figure 1a enumerated realms `S`, `T`, and `W`; the corresponding grammar is shown in Figure 1b. Just as the set of all sentences defines a language, the set of all component compositions defines a *family of systems*. Adding a new component to a realm is equivalent to adding a new rule to a grammar; the family of systems enlarges automatically. Because large families of systems can be built using few components, GenVoca is a *scalable* model of software construction.

Symmetry. Just as recursion is fundamental to grammars, recursion in the form of symmetric components is fundamental to GenVoca. More specifically, a component is *symmetric* if it exports the same interface that it imports (i.e., a symmetric component of realm `W` has at least one parameter of type `W`). Symmetric components have the unusual property that they can be composed in almost arbitrary ways. In realm `W` of Figure 1, components `n` and `m` are symmetric whereas `p` and `q` are not. This means that compositions `n[m[p]]`, `m[n[p]]`, `n[n[p]]`, and `m[m[p]]` are possible, the latter two showing that a component can be composed with itself. In general, the order in which components are composed can significantly

2. Parameterizations that we examine in this paper are simple enough to dispense with formal parameter names.

3. Note that composing components can be interpreted as stacking layers in hierarchical software systems. We use the terms *component* and *layer* interchangeably in this paper.

affect the semantics, performance, and behavior of the resulting system [Bat92a].

Design Rules and Domain Models. In principle, any component of realm `s` can instantiate the parameter of any component of realm `T`. Although the resulting equations would be *type correct,* the equation may not be semantically correct. That is, there are often domain-specific constraints *in addition to* implementing a particular virtual machine that instantiating components must satisfy. These additional constraints are called *design rules*. *Design rule checking (DRC)* is the process of applying design rules to validate type equations.

A *domain model* for a GenVoca generator consists of realms of components and design rules that govern component composition. In the next section, we briefly review the domain model of the P2 generator and illustrate some of its design rules.

3 P2 Domain Model

P2 is a GenVoca generator for container data structures [Bat93-94]. The domain model of P2 relies on two realms: `ds` and `mem`. `ds` components export a standardized container-cursor interface. Among the components of `ds` are those that implement common data structures (e.g., binary trees, doubly-linked ordered and unordered lists) and storage options (e.g., free lists of deleted elements, sequential and random storage). `mem` components export standardized memory allocation and deallocation operations. Among its members are components that manage space in persistent and transient memory.

```
ds = { bintree[ ds ],   // binary tree
       dlist[ ds ],     // unordered doubly
                        // linked list
       odlist[ ds ],    // key-ordered list
       avail[ ds ],     // free list of
                        // deleted elements
       index[ ds, ds ],// key indexing
       malloc[ mem ],  // heap storage
       array[ mem ],   // array storage
       inbetween[ ds ],// deletion actions
       top2ds[ ds ],   // first layer of a
                        // ds expression
       … }

mem = { transient,    // transient memory
        persistent,   // persistent memory
        … }
```

Currently there are over fifty components in P2, most of which are symmetric. Container data structures are defined by type equations that typically reference from five to twenty components. Unfortunately, the correctness of even the simplest equations is not obvious. Validation is complicated by the fact that many components have nonobvious rules for their use.

As an example, the `inbetween` component encapsulates algorithms that are common to many data structure components (e.g., `bintree` and `dlist`). These algorithms deal with the positioning of a cursor immediately after an element has been deleted (e.g., does the cursor point to a "hole" or should it be positioned on the next element in the container?). Instead of replicating these algorithms in every data structure component (and then dealing with the maintenance/consistency problems that would ensue), the algorithms are written once (i.e., factored) as the `inbetween` component. A consequence of this factoring is that a precondition for using a data structure component is the previous appearance of `inbetween` in a type equation. More specifically, the valid use of `inbetween` requires that a *single* copy of `inbetween` be present in a type equation that uses at least one data structure component (`dlist`, `bintree`, etc.) and it should precede *all* such components in the equation. The `right` equation, below, shows a correct usage — i.e., `inbetween` precedes all data structure components. The `wrong` equation, below, shows an incorrect usage: a data structure component `dlist` appears prior to `inbetween`.

```
right = …inbetween[…[dlist[dlist[…]]]]…;
wrong = …dlist[…[inbetween[dlist[…]]]]…;
```

Rules such as this should not be borne by programmers; they are *much* too easy to forget and to be misapplied. A design rule checker that tests such rules automatically and reports errors when they occur removes a great burden from P2 users. We first present a general model of design rule checking in Section 4 and then show how we adapted the model to P2 and Genesis generators in Section 5 and Section 6.

4 A Model of Design Rule Checking

Perry's Inscape is an environment for managing the evolution of software systems [Per87-89b]. Among the features it supports is consistency checking, a simplified form of verification. Components (i.e., operations) have preconditions for their use and postconditions (that describe what is known to be true as a result of an operations's execution). A novel aspect of Inscape is that components additionally have *obligations* which are conditions that must be satisfied by any system that uses a component. Obligation predi-

cates require “action-at-a-distance”: although they *might* be satisfied locally by adjacent components, generally they depend on global properties of the system (i.e., on properties of nonadjacent components). Obligations are propagated to their enclosing modules where eventually they must be satisfied by some postconditions. Another aspect of Inscape is that full-fledged verification is not attempted. Instead, primitive predicates are declared and informally defined, typically with their names hinting at their semantics. Preconditions, postconditions, obligations are expressed in terms of these predicates, thus enabling a practical but powerful form of “shallow” consistency checking to be achieved using pattern matching and simple deductions.

The Inscape approach can be adapted to design rule checking by exploiting the semantics of layers. First, design rule checking examines states of software system (type equation) development; it does *not* model states of system execution. Figure 2 illustrates the distinction. Suppose `s[Q]` is a system that is parameterized by realm `Q`. Suppose further that `k[...]` is a component of `Q`. Composing `s` with `k` maps system `s` to system `s' = s[k[...]]`. To model states of system (type equation) development, every system is described by a set of attributes whose values define its states or properties. Thus, we might define an attribute `State` whose value is `no-loops` in system `s` (meaning that `s` has no loops), and after instantiation, `State` has the value `has-loops` (meaning that `s'` has loops). *Design rule checking deals with the testing and assignment of system design states; it assumes that all transformations (components) are semantically correct.*

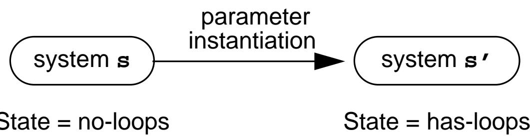

Figure 2. Modeling States of Program Development

Second, it is common for GenVoca components to have preconditions and obligations that are not satisfied locally, i.e., by components that are adjacent to it in a type equation. Preconditions and obligations of a component `k` are satisfied “at-a-distance”, that is, by components that either lie (far) beneath `k` or (far) above `k` in a type equation.[4] Moreover, the properties exported by `k` to “higher” layers are generally *not* the same properties that are exported to “lower” layers. For this reason, we found it necessary to distinguish two kinds of preconditions and postconditions.[5]

Postconditions are properties of `k` that are to be exported to components *beneath* `k` in a type equation. *Preconditions* define the properties that must hold for `k` to work properly; they test the cumulative postconditions of components that lie *above* `k` in a type equation.

> **Example**. Suppose component `k` has a precondition that attribute `A` must have the value `v` (see Figure 3a). For `k` to be used correctly, there must be some component, say `u`, that sits above `k` whose postcondition sets `A = v`. Note that `u` need not be immediately above `k`; `u` might reside far above `k`.

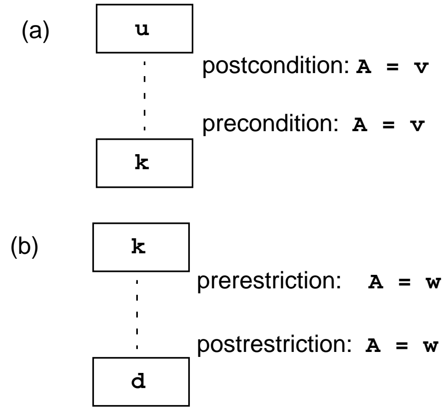

Figure 3. Different Kinds of Design Rules

Postrestrictions are properties of `k` that are to be exported to components *above* `k` in a type equation. *Prerestrictions* (which correspond to Inscape obligations) are preconditions for instantiating component parameters; they test the cumulative postrestrictions of components that lie *beneath* `k` in a type equation.

> **Example**. Suppose component `k` has a single parameter with the prerestriction that attribute `A` must have the value `w` (see Figure 3b). For the parameter to be correctly instantiated, there must be some component, say `d`, that lies below `k` whose postrestriction sets `A = w`. Analogously, `d` need not be immediately beneath `k`; `d` might reside far below `k`.

Given GenVoca design rules (i.e., preconditions, postconditions, prerestrictions, and postrestrictions)

4. We use the terms “higher” and “lower” refer to relative positions of components within a type equation. The outermost component of an equation is the “highest” component, and the innermost components are the “lowest”.

5. There may be some dispute on the proper terminology to use; preconditions and postconditions usually refer to run-time properties, not design-time properties. As there seems to be no commonly used terms for design-time preconditions and postconditions, we chose not to invent more terms.

of every component of a type equation, design rule checking involves:

- a top-down propagation of postconditions and the testing of component preconditions, and
- a bottom-up propagation of postrestrictions and the testing of parameter prerestrictions.

In the following sections, we present general algorithms for top-down and bottom-up design rule checking. We initially place no restrictions on the complexity of DRC predicates. Later in Section 5, however, we show that predicates for domain-customized instances of our algorithms are very simple and are consistent with the shallow consistency checking approach taken in Inscape [Per87-89b].

4.1 Top-Down Design Rule Checking

Consider component `k[x]` which has a single parameter `x`. `k` has both a precondition (`precondition-k`) and a postcondition (`postcondition-kx`). Let `top` denote the set of attribute values that are known to hold at the point immediately above `k` in a type equation. Component `k` is correctly used if `top` implies `k`'s preconditions (i.e., `top` $\Longrightarrow$ `precondition-k`). The set of attribute values that hold immediately beneath `k` in the type equation is computed by applying the postconditions of `k` to the current conditions (i.e., `top-x` = `postcondition-kx` $\oplus$ `top`). The left-associative operator $\oplus$ is the *postcondition propagation operato*r. When type equations correspond to a linear stack of components, the testing of preconditions and the propagation of postconditions is straightforward: only two operators $\oplus$ and $\Longrightarrow$ are needed.

In general, type equations are trees of components. Branching arises when components have multiple parameters, e.g., `d[x,y]`. Each parameter of a component has its own postcondition that defines the set of attribute values that hold for that parameter; these are the values that are propagated to any system instantiating that parameter. In the case of component `d[x,y]`, parameter `x` would have `postcondition-dx` as its postcondition and parameter `y` would have `postcondition-dy`.[6] Let `top` is the set of conditions that hold prior to component `d` in a type equation, `top-x` be the set of conditions that hold for parameter `x` after `d` has been applied, and `top-y` be the set of conditions that hold for parameter `y`. `top-x` is computed by applying `x`'s postcondition to `top` (i.e., `top-x` = `postcondition-dx` $\oplus$ `top`) and `top-y` is computed similarly (`top-y` = `postcondition-dy` $\oplus$ `top`). Given the operators $\oplus$ and $\Longrightarrow$, there is a straightforward, recursive algorithm for the top-down propagation of postconditions and the testing of component preconditions [Bat95].

6. Postconditions for different parameters are generally not the same. The realm of a parameter can be expressed as a postcondition. If a component had two parameters and the realms for both were different, so too would be their postconditions.

4.2 Bottom-Up Design Rule Checking

Every parameter of a component has preconditions (called *prerestrictions*) for instantiation; every component also has postconditions (called *postrestrictions*) that are exported to higher-level layers in a type equation. Figure 4 depicts a typical situation: components `q`, `r`, `s`, `t`, and `w` are composed hierarchically, and `q` has a single parameter. In general, the prerestrictions for `q` are not satisfied by the component `r` that instantiates its parameter, but rather by components deep within the system rooted at `r`. That is, the prerestrictions of `q` may be satisfied by `r` or `s` or `t` or `w`, or any combination thereof.

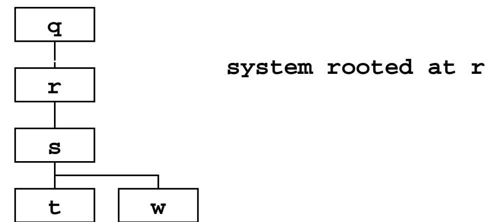

Figure 4. System Instantiation of Parameters

This gives rise to a different interpretation of instantiation, namely that *systems* instantiate parameters, not *components*. Every system exports a realm interface plus a set of attribute values (called *system postrestrictions*) that higher-level layers can reference. A component parameter has been correctly instantiated if the postrestrictions of the instantiating system imply that parameter's prerestrictions.

Consider component `u[x]`. `u` has both a prerestriction (`prerestriction-ux`) and a postrestriction (`postrestriction-u`). Let `bottom` denote the set of attribute values that are exported by a system that instantiates parameter `x`. `x` is instantiated correctly if `bottom` implies its prerestrictions (i.e., `bottom` $\Longrightarrow$ `prerestriction-ux`). The set of attribute values that are exported by the system rooted at `u` is computed by applying the postrestrictions of `u` to the attribute values of the system that it imported (i.e., `bottom'` = `postrestriction-u` $\oplus$ `bottom`). Note that the same operators $\Longrightarrow$ and $\oplus$ used in top-down design rule checking are used in bottom-up design rule checking. Just as in the case of top-down design rule checking, there is a simple, recursive algorithm for the bottom-

up propagation of postrestrictions and the testing of parameter prerestrictions [Bat95].

4.3 Attribute Grammars

McAllester [McA94] observed that attribute grammars unify realms, components, attributes, top-down and bottom-up design rule checking. From previous sections, we know that realms of components define a grammar. Attributes model states of system (type equation) development, where postconditions assign values to inherited attributes (i.e., attributes whose values are determined by component ancestors) and postrestrictions assign values to synthesized attributes (i.e., attributes whose values are determined by component descendants). The practical benefit of this connection with attribute grammars, besides the fact that design rule checking reduces to a well-studied problem, is that common tools, such as **lex** and **yacc**, are well-suited for writing design rule checkers, as we'll see in Section 6.

5 Targeting DRC Algorithms to Specific Domains

The design rule checking algorithms of Section 4 are domain-independent. To specialize them to a particular domain, we need definitions and representations for attributes, predicates, and the operators ⊕ and ⇒. In the following, we explain the representations that we implemented for P2; virtually the same representations were used in Genesis.

5.1 Attributes

An attribute models a property that exposes a composition constraint. Although the properties in which we are interested undoubtedly have complex formal definitions, we have found (like Perry [Per87-89b]) that in practice they can be defined informally as attributes that assume restricted values. The values we use (**any**, **assert**, **negate**, and **inherit**) are defined in Table 1.

Example P2 attributes are: **df_present** and **retrieval**. **df_present** represents the property that a component implements logical deletions. That is, instead of physically deleting an element for a container, the component marks the element deleted but does not immediately reclaim its space. The **retrieval** attribute represents the property that a component interlinks all elements of a container to facilitate searching. Components that implement data structures (e.g., **bintree**, **dlist**, etc.) have the **retrieval** property. The assignment of **assert** or **negate** to these attributes as a postcondition or postrestriction depends on whether a component satisfies the property. **inherit** is used when the value of an attribute is unchanged by a component.

Attribute Value	Interpretation
any	nothing is known
assert	property is asserted
negate	property is negated
inherit	property value is inherited from existing conditions

Table 1. Attribute Values used in P2 and Genesis

5.2 Predicates

Preconditions and prerestrictions in P2 and Genesis request specific attribute values (e.g., **any**, **assert**, **negate**), but not how the attribute value was determined (e.g., **inherit**). Table 2 lists the four different primitive predicates that can be defined over a *single* attribute.[7] P2 predicates are simple conjunctions and disjunctions of primitive predicates. Conjunctive predicates, for example, encoded as a vector of primitive predicates that are indexed by attribute. Thus, predicate $P_1 \wedge P_2 \wedge \ldots \wedge P_n$ would be encoded as the vector $[P_1, P_2, \ldots, P_n]$ where P_i is the primitive predicate for attribute *i*.

Predicate	Interpretation
true (any)	true (no constraints)
assert	attribute has assert value
negate	attribute has negate value
false	false (unsatisfiable)

Table 2. Primitive Predicates used in P2 and Genesis

5.3 Postcondition Propagation Operator ⊕

Component postconditions and postrestrictions selectively declare new attribute values (e.g. **assert** or **negate**) or propagate existing (**inherited**) values. Table 3 defines the condition propagation operator + for a *single* attribute. Given a postcondition/postrestriction value vector $V = [V_1, V_2, \ldots, V_n]$ and the vector of existing conditions $E = [E_1, E_2, \ldots, E_n]$, the ⊕ operator is vector addition (using the + operator of Table 3):

$$V \oplus E = [\, V_1 + E_1,\ V_2 + E_2,\ \ldots,\ V_n + E_n \,]$$

7. In [Bat95] we explain the need for a fifth attribute value. This enlarges the set of primitive predicates.

Postcondition/Postrestriction + Existing Condition		Existing Condition			
		true	**assert**	**negate**	**false**
Postcondition or Postrestriction	**assert**	assert	assert	assert	assert
	negate	negate	negate	negate	negate
	inherit	true	assert	negate	false

Table 3. The Propagation Operator + for a Single Attribute

Existing Condition → Precondition/ Prerestriction		Precondition or Prerestriction			
		true	**assert**	**negate**	**false**
Existing Condition	**true**	true	false	false	false
	assert	true	true	false	false
	negate	true	false	true	false
	false	true	true	true	true

Table 4. The Implication Operator → for a Single Attribute

5.4 Implication Operator $\Longrightarrow$

The implication operator $\longrightarrow$ for a *single* attribute is defined by a truth-table (Table 4). Given a vector of existing conditions $E = [E_1, E_2, \ldots, E_n]$ and a precondition/prerestriction vector $P = [P_1, P_2, \ldots, P_n]$ (of a conjunctive predicate) the implication operator $\Longrightarrow$ has a simple definition: all primitive predicates must be true for the compound predicate to be true. (A simple generalization handles disjunctions).

$$E \Longrightarrow P = (E_1 \rightarrow P_1) \wedge (E_2 \rightarrow P_2) \wedge \ldots \wedge (E_n \rightarrow P_n)$$

6 Implementation Notes

The implementation of our DRC algorithms and the P2/Genesis specializations of the ⊕, $\Longrightarrow$, and Δ operators was straightforward: the source files consist of 1500 lines of **`lex`** and **`yacc`**. We wrote a general utility, called **`dreck`**, that would allow designers to declare realms, components, and their design rules based on the representations we noted previously for attributes, predicates, and DRC operators [Bat95]. Figure 5 shows a **`dreck`** declaration of the **`array`** component and its design rules. A component's name, realm membership, and realm parameters are declared on the first line. Subsequent lines define design rules. A precondition for **`array`**'s usage is that a layer above **`array`** needs to support logical deletion. This precondition is expressed by asserting the **`df_present`** property. Another design rule is to assert to layers above and below that **`array`** is a retrieval layer. Such a declaration is expressed by asserting the **`retrieval`** property as a postcondition and postrestriction.

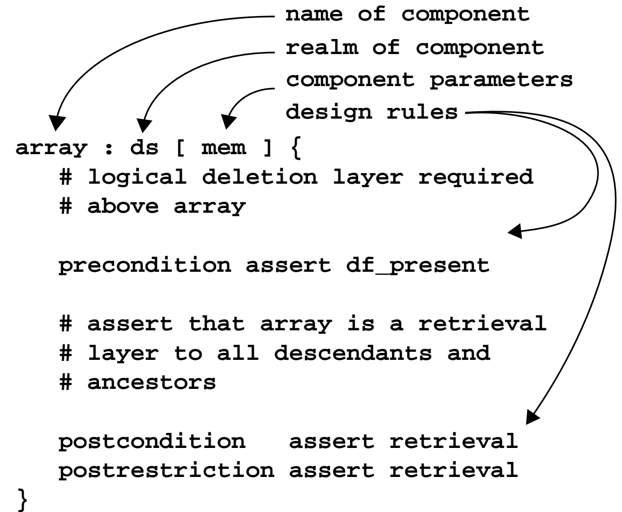

Figure 5. Specification of Design Rules

Algorithm Efficiency. Let n denote the number of components in a type equation and let m denote the number of attributes. A straightforward implementation of the DRC algorithms is as a tree traversal, where each node is visited twice (once on the way down from the root, and once on the way up from visiting leaves). At each visit, m attribute values are propagated. Thus, the complexity of our $O(mn)$.

To give readers upper estimates of n and m, the most complicated type equations that we have encountered in Genesis and P2 have approximately 30 components (i.e., $n \leq 30$). Genesis maintains the greatest number of attributes (m=14), whereas P2 has fewer (m=8), even though both generators maintain a library of 50 components. Although it is not difficult

to envision greater values for m and n, substantially greater values (e.g., $m, n > 100$) seem unlikely.

Extensibility. Adding new components to a domain model is not difficult. The component designer must determine whether existing attributes are sufficient to capture illegal compositions (in which case component addition is trivial) or whether new properties need to be added. In practice, adding more attributes has not been problematic because the number of components in generator libraries is modest (and because of scalability, we would expect the number to remain small). For example, ADAGE has the largest library (about 400 components) which avionics experts have no difficulty managing.

Explanation-Based Error Reporting. Detecting precondition and prerestriction errors is only part of the problem of debugging type equations; repairing equations are also important. Precondition ceilings is a technique used in Inscape that we found particularly effective. Suppose component `Y`'s precondition `A=v` failed. This means that some component above `Y`, say `X`, set $A \neq v$ as a postcondition. To repair this error, there needs to be another component, `Z`, that must be inserted below `X` and above `Y` whose postcondition is `A=v`. Techniques such as this (including obligation/prerestriction ceilings) form the basis of a powerful explanation-based error reporting scheme. The following example illustrates the idea.

Example. Suppose we would like a P2 container implementation that stores elements in a binary tree, whose nodes are stored sequentially in transient memory. A first attempt at a composition might be:[8]

```
first_try =
      top2ds[bintree[array[transient]]];
```

Our DRC algorithms report the following:

```
Precondition errors:
   an inbetween layer is expected between
     top2ds and bintree
   a logical deletion layer is expected
     between top2ds and array

Prerestriction error:
   parameter 1 of top2ds expects a
     subsystem with a qualification layer
```

The first error reminds us (from Section 3) that we forgot that a `bintree` layer requires the `inbetween` layer to be above it. Not only that, the error message states exactly how to repair the equation; there is only one location where `inbetween` can go (i.e., in between `top2ds` and `bintree`). The second error reminds us that `array` requires a logical deletion layer above it. Further, this layer must be below `top2ds`. The third error tells us that a qualification layer is required below `top2ds`. Users with minimal experience with P2 are able to repair all of these errors easily. But suppose repairs lead to the following equation:

```
second_try = top2ds[inbetween[bintree[
   qualify[delflag[array[transient]]]]]];
```

where `qualify` is a qualification layer and `delflag` is a logical deletion layer. The DRC response to this equation is:

```
Precondition error:
   a retrieval layer (bintree) is not
     expected above qualify
```

This error tells us that all retrieval layers must lie beneath `qualify`; the fix is to transpose `bintree` and `qualify`, which results in a correct equation:

```
correct = top2ds[inbetween[qualify[
   bintree[delflag[array[transient]]]]]];
```

In general, DRC error messages direct users to modify an incorrect equation to the nearest set of correct type equations in the space of all equations. We have found this advice works well. With minimal experience, P2 users typically come very close to their desired equation on the first attempt; DRC messages enable them to correct errors quickly.

7 Related Work and Insights

Related Work. DRACO used a form of shallow consistency checking (called assertions and conditions) in composing layers of transformations [Nei80]. DaTE, the design rule checker for Genesis [Bat92b] supported only component preconditions. The limitations of DaTE led to the work presented in this paper.

McAllester developed a functional programming language, VAG, based on variational attribute grammars, to address the design rule checking issues for the ADAGE generator [McA94]. Preconditions and prerestrictions are treated uniformly as constraints. The constraints associated with a component are expressed as a VAG program. When an avionics system is composed from components, the set of constraints that must be satisfied is defined by the composition of corresponding VAG programs. The

8. `bintree` links elements of a container onto a binary tree; the nodes of the binary tree will be stored sequentially in an `array`; the array will reside in `transient` memory. The `top2ds` layer must root all P2 type equations; had `top2ds` been absent, the DRC algorithms would report additional errors.

VAG interpreter has limited reasoning abilities to infer values of unbound VAG program parameters.

Parameterized programming is intimately associated with the verification of component compositions. Goguen's work on OBJ and library interconnection languages, such as LIL and LILEANNA [Gog86, Tra93], are basic. The RESOLVE project explores the design of reusable and parameterized components, component certifiability, and the certifiability of component compositions [Sit94]. Although there are many similarities among these works and ours, there is a basic difference: there is no "action-at-a-distance" in the other work. Compositions of OBJ, LILEANNA, and RESOLVE components are verified locally; components constrain the behavior of immediately adjacent components, and not components that reside far above or below them in a hierarchy.

Our work is also an example of the types of consistency checking problems encountered in software architectures [Per92, Gar94-95, Mor94]. To our knowledge, other than Inscape, validating compositions of components in the context of architectures has only begun to be addressed.

Insights. Our work on DRC was actually developed independently of DRACO and Inscape. That our results are so similar is encouraging: we suspect that "shallow" consistency checking is a general technique for automatic software system generation.

An important distinction between Inscape and our work is the scale of componentry. An Inscape component is a function; a GenVoca component is a subsystem (i.e., a suite of interrelated classes). Perry noted that there can be many primitive predicates when there are thousands or tens of thousands of functions in a system. In contrast, type equations rarely reference more than fifty components, and the number of primitive predicates that we have encountered in modeling different and multiple domains is modest (i.e., about 10). So, it would seem that scaling the size of a component *reduces* the number of primitive predicates (attributes) that need to be maintained. This seems counterintuitive.

Our best explanation for this centers on two observations. First, we believe that modeling states of software system development (instead of states of execution) reduces the number of properties to examine. Second, we believe that GenVoca offers a powerful methodology for the design of reusable components. Object-oriented design methodologies, for example, are powerful because of their ability to manage and control software complexity [Boo91]. It is not difficult to recognize that standardizing domain abstractions and their programming interfaces (i.e., the core of GenVoca) is also a powerful way of managing and controlling the complexity of software in a *family of systems*. We believe that standardization makes some problems tractable that would otherwise be very difficult. Composibility of software components is one example (c.f., [Gar95]) and DRC is another (c.f. [Kat92]).

8 Conclusions

Software system generators are becoming important tools for software developers. Generators utilize libraries of reusable components to assemble complex, high-performance systems quickly and cheaply. Each library component has limitations, called design rules, on how it can be combined with other components. Experience has shown that validating component compositions is difficult to do by casual inspection; as the number of components and the complexity of their rules grow, a mechanical approach to validation is absolutely essential.

We have shown that a GenVoca domain model is an attribute grammar, where sentences of the grammar define valid compositions of components. We have shown how the shallow consistency checking approach of Perry's Inscape environment can be adapted to exploit the semantics of GenVoca layers to define the actions of GenVoca production rules. Our approach distinguishes predicates and properties of component usage from those of parameter instantiation. We have shown (and experience confirms) that domain-specific instances of our algorithms are practical: they are simple, easy to implement, and efficient. Moreover, they offer powerful explanation-based error reporting capabilities to suggest how incorrect compositions can be repaired.

Finally, we have observed that the number of attributes (primitive predicates) that need to be maintained for design rule checking GenVoca components is rather small. This is in contrast to small-scale components (i.e., functions) where the number of primitive predicates to be maintained can become large. We believe the explanation for this lies in the power of standardization to control the complexity of families of software systems. Components that are designed to be interoperable, plug-compatible, and interchangeable often make otherwise difficult problems tractable.

So that others may learn from our work, `dreck` is available free of charge via the Predator web page: `http://www.cs.utexas.edu/users/schwartz/.`

Acknowledgments. We thank Dewayne Perry for stimulating discussions on Inscape's and drafts of this paper. We also thank Ira Baxter, Paul Clements, Dave Weiss, Chris Lengauer, Bruce Weide, and Steve Edwards for their insights on this work. Finally, we are grateful to Mike Hewett and Chris Petrock for their CS395T term project that gave rise to this paper.

9 References

[Bat92a] D. Batory and S. O'Malley, "The Design and Implementation of Hierarchical Software Systems with Reusable Components", *ACM TOSEM*, October 1992.

[Bat92b] D. S. Batory and J. R. Barnett. "DaTE: The Genesis DBMS Software Layout Editor." In *Conceptual Modeling, Databases, and CASE*, Pericles Loucopoulos and Roberto Zicari, eds. John Wiley & Sons, 1992.

[Bat93] D. Batory, et al., "Scalable Software Libraries", *Proc. ACM SIGSOFT,* December 1993.

[Bat95] D. Batory and B.J. Geraci, "Validating Component Compositions in Software System Generators", UT/CS TR-95-03, University of Texas at Austin, 1995.

[Bax92] I. Baxter, "Design Maintenance Systems", CACM April 1992, 73-89.

[Bla91] L. Blaine and A. Goldberg, "DTRE - A Semi-Automatic Transformation System", in *Constructing Programs from Specifications*, Elsevier Science Publishers, 1991.

[Boo91] G. Booch. *Object-Oriented Design With Applications*, Benjamin-Cummings, 1991.

[Big94] T. Biggerstaff, "The Library Scaling Problem and the Limits of Concrete Component Reuse", *ICSR* 1994.

[Coh95] S. Cohen, et al., "Models for Domains and Architectures: A Prescription for Systematic Software Reuse", *AIAA Computing in Aerospace*, 1995.

[Cog93] L. Coglianese and R. Szymanski, "DSSA-ADAGE: An Environment for Architecture-based Avionics Development", *Proc. AGARD*, 1993.

[Gar94] D. Garlan, et al., "Exploiting Style in Architectural Design Environments", *ACM SIGSOFT 1994.*

[Gar95] D. Garlan, et al, "Architectural Mismatch or Why It's Hard to Build Systems out of Existing Parts", *ICSE 1995*.

[Gog86] J.A. Goguen, "Reusing and Interconnecting Software Components", *Computer*. February 1986.

[Gom94] H. Gomaa, et al., "A Prototype Domain Modeling Environment for reusable Software Architectures", *ICSR* 1994.

[Gri94] M.L. Griss and K.D. Wentzel, "Hybrid Domain-Specific Kits for a Flexible Software Factory", *ACM SAC'94*, March 1994.

[Hei93] J. Heidemann and G. Popek, "File System Development with Stackable Layers", *ACM TOCS*, March 1993.

[Hut91] N. Hutchinson and L. Peterson, "The *x*-kernel: An Architecture for Implementing Network Protocols", *IEEE TSE*, January 1991.

[Joh92] R.E. Johnson, "Documenting Frameworks using Patterns", *OOPSLA 1992*, 63-76.

[Kat92] M.D. Katz and D.J. Volper, "Constraint Propagation in Software Libraries of Transformation Systems", *Int. Jour. Software Engineering and Knowledge Engineering*, Vol. 2#3 (1992).

[Lei94] J.C.S. do Prado Leite, et al., "Draco-PUC: A Technology Assembly for Domain-Oriented Software Development", *ICSR* 1994.

[McA94] D. McAllester, "Variational Attribute Grammars for Computer Aided Design." ADAGE-MIT-94-01.

[Mor94] M. Moriconi and X. Qian, "Correctness and Composition of Software Architectures", *ACM SIGSOFT 1994.*

[Nei80] J. Neighbors, "Software Construction Using Components", Ph.D. Thesis, ICS-TR-160, University of California at Irvine, 1980.

[Nin94] J.Q. Ning, et al. "An Architecture-Driven, Business-Specific, and Component-Based Approach to Software Engineering", *ICSR* 1994.

[Per87] D.E. Perry, "Software Interconnection Models", *ICSE 1987*.

[Per89a] D.E. Perry, "The Logic of Propagation in The Inscape Environment", *ACM SIGSOFT 1989*.

[Per89b] D. E. Perry, "The Inscape Environment", *Proc. ICSE 1989*.

[Per92] D.E. Perry and A.L. Wolf, "Foundations for the Study of Software Architecture", *ACM SIGSOFT Software Engineering Notes*, October 1992.

[Sit94] M. Sitaraman and B. Weide, "Component-Based Software using RESOLVE", *ACM Software Engineering Notes*, October, 1994.

[Tra93] W. Tracz, "LILEANNA: A Parameterized Programming Language," *ICSR* 1993.

[Ude94] J. Udell, "Componentware", *BYTE*, May 1994.

[Wei90] D.M. Weiss, *Synthesis Operational Scenarios*, Technical Report 90038-N. Version 1.00.01, Software Productivity Consortium. August 1990.