

# Demeter Interfaces: Adaptive Programming without Surprises

Therapon Skotiniotis

Jeffrey Palm

Karl Lieberherr

College of Computer & Information Science  
Northeastern University  
360 Huntington Avenue  
Boston, Massachusetts 02115 USA  
{skothe,jpalm,lieber}@ccs.neu.edu

## ABSTRACT

Adaptive Programming (AP) allows for the separate definition of data structures, traversals and computations over these data structures. The loosely defined contexts, structure and computation, are composed according to a given traversal specification. Traversal specifications are defined against a graph-based model of the underlying data structure with the ability to abstract over graph node names, edges and subpaths. As such certain alterations to the data structure will not affect the meaning of the program. Currently AP systems, *e.g.* DemeterJ, provide no mechanisms to warn or even guard against modifications that *will* affect the meaning of a program. Programmers have to depend on thorough testing in order to detect such modifications.

In this paper we present *Demeter Interfaces*, an extension to the DemeterJ system. Demeter Interfaces specify the expected structural properties of the underlying data structure that must hold in order for adaptive code to function correctly. Through an example we illustrate the usage and implementation of Demeter Interfaces. We further illustrate how Demeter Interfaces result in better designs of adaptive programs, ease of adaptive code reuse and how they promote parallel development. We discuss how Demeter Interfaces are applicable to XML programs that use the XPath navigation operator // . This makes our results relevant to large communities, such as the XQuery and XSLT communities.

## 1. INTRODUCTION

Adaptive Programming [4](AP) takes the principle of Structural Recursion (SR) to a new level. In SR you organize your program following the structure of your data, a design principle used today for introducing program design [?] . AP applies SR by focusing on the difference between the structural recursion template and the desired program. The information in the structural recursion template (that is systematically defined by the data type definition), is not repeated in the adaptive program.

There are many ways of defining a Structural Recursion template. There is a maximal one that uses the entire data type. But of-

ten only a subset of the data type is relevant. AP introduces traversal strategies to define the relevant Structural Recursion template. Traversal strategies should be viewed as a first step to customizing Structural Recursion templates. There are other ways to parameterize the Structural Recursion template: is it functional or imperative; which information should be carried down the structure and how far.

An adaptive program is a program written in terms of loosely coupled contexts, *i.e.*, data structure and behavior (computations) with a third definition succinctly binding the two contexts together. In DemeterJ [6], the most popular incarnation of AP, programs are defined in terms of their data structures and a set of computations that take place during the traversal of data structures. Structure and computation come together through *traversal specifications*. Traversal specifications can abstract over graph node names, edges and subpaths that allows certain modifications to the underlying data structure without altering the meaning of the whole computation. Traversal specifications are defined using a domain specific language that operates on a graph-based model of a program's data structure.

Using Aspect Oriented Programming (AOP), AP expresses the difference between the recursion template and the desired program. AP introduced a dynamic join point model based on object entry and exit events. AP had as key motivation to incrementally enhance a Structural Recursion template into the desired program.

Identifying which modifications do not alter the meaning and intent of the computation is not a trivial task. Identifying which parts of the data structure are explicitly referenced in traversal specification and their attached computations is not enough. Assumptions made by the adaptive code about the underlying data structure have to be also verified. As a simple example consider a traversal specification used to obtain an object deep inside a data structure, *e.g.*, “*from Company to CEO*”, on a data structure representing a company. If this traversal specification was written with the assumption that there is only **one** CEO in a company and a modification is made to allow multiple CEOs then no compile time error is issued by DemeterJ. The program compiles and executes, however the result of the computation will be incorrect. DemeterJ offers no way to define and check for such assumptions. Programmers have to rely on extensive tests that need to be run everytime a modification is made, defeating the benefits of adaptive programming.

The problem of modifications that alter the meaning of the program makes iterative and parallel development difficult. As dependencies between computations and traversals arise, it becomes harder to properly test and detect bugs in adaptive programs. At the same time, as AP programs become larger in size understandability

of adaptive code decreases. Programmers have to both look at the traversal specification and the computation attached to it but also calculate from the concrete data structure the possible paths based on the traversal specification. Finally, since both traversal specifications and their attached computations refer to concrete data structure names, decrease reusability of adaptive code.

In this paper we introduce Demeter interfaces that make AP without surprises. The Structural Recursion template is defined by a high-level traversal strategy on a data type, and the program is defined as an enhancement to the template. If the data type changes, the program might still work without any changes to the strategy nor the aspects that enhance the template. This gives adaptive programs an impressive flexibility to data type changes. But this flexibility is also dangerous: For the new data type, the program might work but give the wrong result. Demeter interfaces are intended to restrict the set of data types that are allowed to be used with an adaptive program.

Demeter Interfaces hit a sweet spot between flexibility and safety. They restrict a bit what AP can do but without going back to the old way of writing the Structural Recursion template manually. They are safer because there is now more intent that is used to check any future data type against an adaptive program. As a side effect, adaptive programs become better documented and more understandable.

In this paper we introduce Demeter Interfaces (DIs) as a mechanism within DemeterJ that allows programmers to define:

- traversal strategies,
- an *Interface Class Graph (icg)*, which provides an interface for the concrete data structure
- a *Visitor Interface*, used to implement the attached computation
- a set of *Constraints*, that define properties that both the icg and the underlying data structure must satisfy

The concrete data structure definition is extended to accommodate for a name mapping between its data members and those of the DI that it implements. The DemeterJ system is extended to statically verify the mapping and validate all DI constraints.

ADD PARAGRAPH FOR DESCRIPTION OF EACH SECTION.

## 2. DEMETER INTERFACES

Demeter Interfaces provide a view of an adaptive program that encloses all the necessary assumptions and structure properties of the underlying concrete data structure, traversal specifications and a visitor interface. Computation is attached to the traversal via the implementation of the visitor interface.

The rest of this section illustrates the usage of DI and their advantages through an equation system example and the implementation of a semantic checker. We first provide a solution in DemeterJ [6] describing DemeterJ. We then iteratively extend the equation system, exposing the issues with the current DemeterJ implementation. We then show a solution that uses DIs and analyzing the advantages over our initial implementation.

### 2.1 A simple equation system in DemeterJ

Our example is about systems of equations in which we want to check that all used variables are defined (we call this a *semantic checker*). A simple equation system could be  $x = 5; y = 9; z = x + y;$ .

Adaptive programs in DemeterJ are defined through a Class Dictionary (**cd**), a set of Behavior Files (**beh**) and Java code. Listing 1 and Listing 3 make up the DemeterJ code for the simple equation system and the implementation of the semantic checker which prints out two lists; defined variables and used variables.

Listing 1: Class Dictionary for Simple Equations

---

```
EqSystem = <equations> BList(Equation).
Equation = <lhs> Variable "=" <rhs> Expr.
Expr : Simple | Compound .
Simple : Variable | Numerical.
Variable = Ident.
Numerical = <val> Integer.
Compound = "(" <lrand> Expr <op> Op <rrand> Expr)".
Op : Add | Sub | Mul | Div.
Add = "+".
Sub = "-".
Mul = "*".
Div = "/".
BList(S) ~ "(" S {";" S} ")".
visitor DefCollector = .
```

---

A cd file is a textual representation of the OO structure of the program which specifies classes and their members. Figure 1 shows the UML representation of the class dictionary in Listing 1. Each line of the class dictionary defines a class definition. An equal sign (“=”) defines a concrete class with the class name on the left hand side of the equals and the members of the class in the right hand side of the equals. Replacing the equal sign with a colon (“:”) defines an abstract class with its subclasses on the left of the colon. Text within left and right angle braces defines class member variable names. The DemeterJ system uses a class dictionary as a grammar definition, providing a language that can parse in definitions and create the appropriate objects. Tokens in the class dictionary surrounded in quotes define the generated languages’ syntax tokens (Listing 2). Parameterized classes are defined through the tilde (“~”) operator. The **visitor** specifies a visitor class.<sup>1</sup>

Listing 2: An instance of a simple equations system given as input to DemeterJ

---

```
(
  x = 5;
  y = (x - 2);
  z = ((y-x) + (y+9))
)
```

---

Behavior files are used to introduce code to the classes specified in the accompanying class dictionary. Introductions can be in the form of regular Java methods, adaptive methods or visitor methods. In Listing 3 the definition of `main` is introduced into the class `Main` as a regular Java method.

Listing 3: Behavior File for Simple Equations

---

```
Main {
  public static void main(String args[]) throws Exception { {
    EqSystem eqSystem = EqSystem.parse(System.in);
    System.out.println("Input is : ");
    eqSystem.print();
    System.out.println("");
    System.out.println("Defined are : ");
    eqSystem.printDefined();
    System.out.println("");

    System.out.println("Used are : ");
```

---

<sup>1</sup>The `DefCollector` class has no data members.

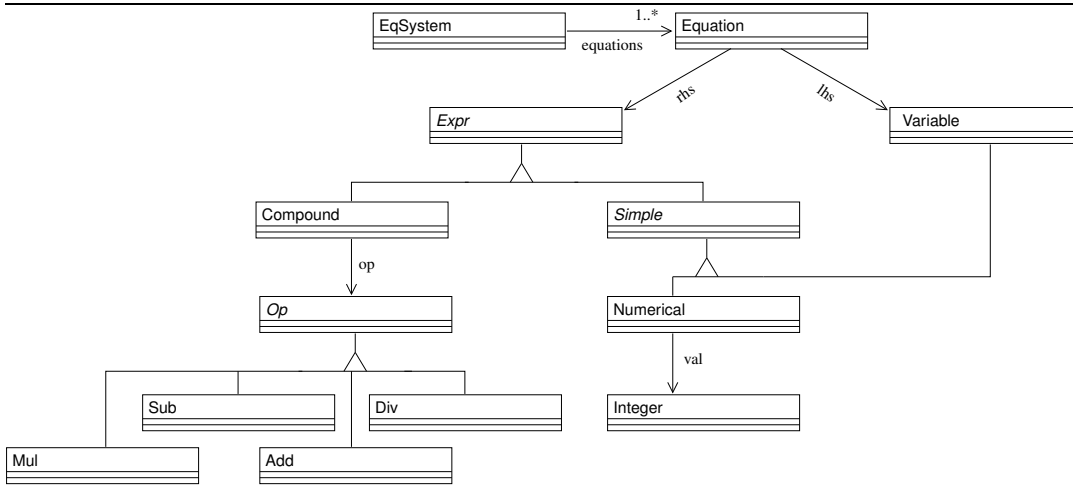


Figure 1: The UML equivalent of Simple Equations Class Graph.

```

    eqSystem.printUsed();
    System.out.println("");
  }}
}

EqSystem {
  public void print() to * (PrintVisitor);

  public void printDefined() bypassing -> *,rhs,* to Ident (
    DefCollector);

  public void printUsed() bypassing -> *,lhs,* to Ident (
    DefCollector);
}

DefCollector {
  before Variable {{
    System.out.println("Found Ident : " + host.get_ident());
  }}
}

```

Three adaptive method signatures are introduced into EqSystem, called `print`, `printDefined` and `printUsed`. Adaptive method signatures are similar to normal Java method signatures extended with a traversal specification, e.g., to `*`, and passed the name of a Visitor class (The `PrintVisitor` prints all instances in a traversal and it is automatically generated by the system). The adaptive method `printUsed` will introduce a method into `EqSystem` that takes no arguments and returns void. When called, the method will traverse from an instance of `EqSystem` to instances of `Ident`. The traversal is not allowed to go through any edge with the label `lhs`. At every instance that the traversal comes across the `DefCollector` visitor definition is advised. If there is a visitor directive, e.g., **before**, **after** or **around**, and the directives argument matches the instance type, the code attached to that directive is executed. Before traversing a `Variable` instance the traversal will print out its contents.<sup>2</sup>

With the completed AP implementation of the semantic checker in place we can now evaluate our solution and verify the claims made, both in favour and against, AP. The principle behind AP states [5]

“A program should be designed so that the interface

<sup>2</sup>`Ident` is a DemeterJ wrapper class for Java’s `String` class.

of objects can be changed within certain constraints without affecting the program *at all*.”

For example, altering the equation system so that equations are now in prefix notation requires a single modification to `Compound`’s definition in the class dictionary:

```

Compound = ``('' <op> Op <lrand> Expr
             <rrand> Exp ``)``.

```

However any alteration that either

- modifies class and/or class member variable names that are explicitly referenced by traversal and/or visitors
- or, breaks an assumption about the system on which an adaptive method depends

alters the program’s behavior. For example, altering the equations system to allow for local variable definitions via let-bindings. This modification breaks two assumptions

1. All variable definitions are found on the left handside of the equals
2. All variables have global scope and thus can be used anywhere.

The adaptive methods, as well as the visitor, depend on these assumptions. However these assumptions are not explicitly captured in AP programs. There is no tool support to stop such modifications. In fact naively extending the equation system to accommodate for functions parameters (Listing 4) will generate a valid AP program with the wrong behavior for the semantic checker.

Listing 4: Class Graph for equations systems with functions of one argument.

```

EquationSystem = <equations> BracketList(Equation).
Equation = <lhs> VarOrFunc "=" <rhs> Expression.
VarOrFunc : DVariable | Function.
Function = "fun" Variable "(" CommaList(Variable) ")".
Expression : *lookahead*(@2@) FunCall | Simple | Compound .
Simple : Variable | Numerical.
Variable = Ident.
DVariable = Ident.
FunCall = Variable "(" CommaList(Simple) ")".

```

```

Numerical = <val> Integer.
Compound = "(" <lrand> Expression <op> Op <rrend>
           Expression)".
Op : Add | Sub | Mul | Div.
Add = "+".
Sub = "-".
Mul = "*".
Div = "/".

```

```

BracketList(S) ~ "(" S { ";" S })".
CommaList(S) ~ S { "," S }.

```

With larger AP programs, it becomes impossible to find all these implicit assumptions and even harder predict which modifications will cause erroneous behavior. Programmers have to rely on exhaustive testing in order to increase their confidence (but not prove) that the program still behaves according to its specification.

To address these problems there needs to be a mechanism that allows for

- the definition of assumptions about the underlying data structure,
- tool support to allow the validation of these assumptions,
- decreased the dependency on class and class member variable names,
- and the modularization of the relevant program information for each adaptive behavior instead of the whole class graph.

## 2.2 A simple equation system with Demeter Interfaces

A Demeter Interface resides between a class graph and the *implementation* of adaptive behavior, i.e., adaptive methods and visitor implementations. A DI is made up of

- an interface class graph [?] (**icg**) that defines a portion of a program structure to which adaptive behavior will be attached,
- traversal strategies that act on the icg,
- constraints imposed on the icg and the traversal strategies,
- adaptive method signatures,
- and a visitor interface.

Figure 2 shows the Demeter Interface for the simple equation system along with a diagrammatical representation (in UML) of the DI's interface class graph.

The **icg** serves as an abstraction of any class graph implementation of the ExprICG DI in order for the strategies defined in ExprICG to be applicable. Furthermore, the constraints defined in the Demeter Interface impose further requirements that need to be met by any concrete implementation of this Demeter Interface. The first constraint expresses that each variable definition introduces exactly one name. The non-empty constraints state that there must be at least one path in the class graph that satisfies the strategy given as parameter. The predicates `unique` and `nonempty` are provided by DemeterJ. The **visitor** section is a list of Java interface definitions for the visitor that need to be implemented for this DI. Method names `before`, `after` and `around` will have the same meaning as the visitor directives in DemeterJ. The **adaptive methods** section holds a list of adaptive method introductions. Adaptive methods now take the visitor as an argument to the method.

Figure 3 gives an example implementation of the ExprICG Demeter Interface (on the right) along with a visitor implementation and a driver class (on the left). `InfixEQSystem` defines the structural relations between concrete classes.

DemeterJ's class dictionaries are extended in two ways; a header is introduced allowing for an **implements** statement that specifies which DIs are being implemented and a mapping between the concrete class dictionary classes to the icg's classes. The concrete class dictionary `InfixEQSystem` provides a definition of its equation system and a mapping  $M$  between the entities in its class graph and the entities in the interface class graph. The mapping definition can map class(es) to class(es), class member variable name(s) to class member variable name(s), strategy to strategy and class to strategy. Invoking the adaptive method requires that a concrete visitor that implements `DVisitor` is given as an argument. The addition of DI requires modifications to the compile time behavior of DemeterJ. For each concrete class dictionary that implements a Demeter Interface

1. Use the DI's mapping  $M$  to ensure that the concrete class dictionary (`InfixEQSystem`) satisfies the interface class graph from the DI (ExprICG), i.e., for each edge  $e$  connecting two nodes  $n_1$  and  $n_2$  in the interface class graph the mapped edge  $M(e)$  connects the mapped nodes  $M(n_1)$  and  $M(n_2)$  in the class dictionary. The tool does support some limited form of inference. After reading in the explicit mapping definitions, the tool follows the structure of the icg and the concrete class graph and attempts to map the remaining entities. The mapping occurs by position for concrete classes and by expanding the definition of abstract classes. If a mapping is not possible an error is reported. For example, in the case of `InfixEQSystem`, `EquationSystem` is mapped to `ESystem` that also maps `Definition` to `Equation`. For the mapping of `Body` the system expands the choices for `Expression` which leads to an element (`Variable` or `Numerical`) or a list of `Expressions` and `Op`. The system takes the alternatives of an `Expression` class and maps those to `Body`.
2. Strategies defined in the DI will be expanded according to the mapping and made available through `InfixEQSystem` class.
3. Each of the constraints will be validated on the expanded strategies for `InfixEQSystem`.

If any of the above steps fails then the concrete class dictionary failed to implement its DI interface. With the simple equation system implemented using Demeter Interfaces we now extend the system and verify that DIs assist the prevention of modifications that alter the program's behavior.

As a first evolution step we want to change from infix notation to prefix notation. This is a modification that does not alter the program's behavior even in the original DemeterJ solution without DIs. Moving to a prefix notation requires to change the definition of `Compound` in `InfixEQSystem` to

```

Compound = <op> Op <lrand> List(Expression)
           <rrend> List(Expression)

```

This change does not affect the Demeter Interface at all. We update the equation system class graph and keep the old mapping. All constraints of the DI are still satisfied after they are mapped into the actual interface class graph and the adaptive methods will function correctly.

```

di ExprICG {
  icg:
    ESystem = List(Definition).
    Definition = <def> DThing "=" <body> Body.
    Body = List(UThing).
    UThing = Thing.
    DThing = Thing.
    Thing = Ident.
    List(S) ~ "(" S ") ".

  strategies:
    gdefinedIdents = "from ESystem via DThing to Ident".
    gusedIdents = "from ESystem via UThing to Ident".
    definedIdent = "from Definition via DThing to Ident".

  constraints:
    unique(definedIdent).
    nonempty(gusedIdents).
    nonempty(gdefinedIdents).

  visitors:
    interface DVisitor {
      public void before (Thing t);
    }

  adaptive methods:
    ESystem {
      public void printDefined(DVisitor) gdefinedIdents;
      public void printUsed(DVisitor) gusedIdents;
    }
}

```

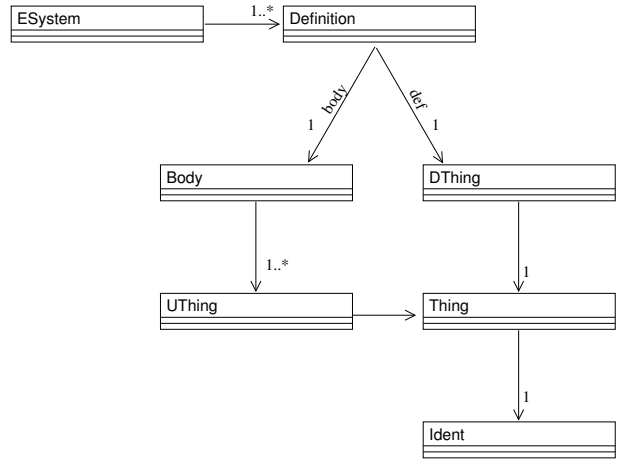


Figure 2: The full definition of the Demeter Interface includes an interface class graph, the strategies on the interface class graph and constraints on the strategies. The UML diagram is equivalent to the interface class graph defined in ExprICG.

It is important to note that during this evolution step, only the DI and the concrete implementation of the interface class graph was needed. Furthermore, the static assurances provided by the tool because of the DI, did not require re-testing of the driver code to ensure that the strategies pick the correct traversals and that the semantic checker still operates as expected. The Demeter Interface allows in this case for separate development and ease of evolution. Hiding irrelevant information through the Demeter Interface provides for higher system modularity.

In the next evolution step we want to add functions with a single parameter to the equation system. This evolution step affects information that is relevant to the semantic checker. The semantic checker has to also deal with parameter names on each function definition but also usages of function definitions that may appear on the right-hand side of equations. Function parameters do not have global scope, their definition scope is local to the function definition. A naive approach would be to alter the concrete class dictionary definition to accommodate for function definitions to appear on the right hand side of equations, i.e.,

Listing 5: Modifications to the concrete class dictionary to accommodate single argument functions

```

cd ParamEquations implements ExprICG{
  EquationSystem = <equations> BracketList(Equation).
  Equation = <lhs> VarOrFunc "=" <rhs> Expression.
  VarOrFunc : DVariable | Function.
  Function = "fun" Variable "(" CommaList(Variable) ")".
  Expression : *lookahead*(@2@) FunCall | Simple | Compound .
  Simple : Variable | Numerical.

```

```

Variable = Ident.
DVariable = Ident.
FunCall = Variable "(" CommaList(Simple) ")".
Numerical = <val> Integer.
Compound = "(" <lrnd> Expression <op> Op <rrand>
  Expression ")".
Op : Add | Sub | Mul | Div.
Add = "+".
Sub = "-".
Mul = "*".
Div = "/".

```

```

BracketList(S) ~ "(" S {"," S } ")".
CommaList(S) ~ S {"," S}.

```

```

use EquationSystem as ESystem,
  (->*,lhs,* to Variable) as DThing,
  (->*,rhs,* to Variable) as UThing,
  Variable as Thing.

```

}

Without altering the mapping and leaving the Demeter Interface intact this approach will result in a compile time error. The mapping definition does not raise an error since DThing can be mapped to Variable by going either through DExpression or through ParametricVariable. The predicate unique for the traversal specification of definedIdent, however, can no longer be satisfied. The modification to allow functions with one parameter breaks one of the assumptions of the interface in particular the

<pre> class Main {     public static void main(String[] args){         DVisitor v = new IDPrintVisitor();         ParamEQSystem ieqs = InfixEQSystem.parse(System.in);         System.out.println("IDs in def:");         ieqs.printDefined(v);         System.out.println("IDs in use:");         ieqs.printUsed(v);     } }  class IDPrintVisitor extends DVisitor {     public void before(Thing id) {         System.out.println(id.toString());     } } </pre>	<pre> cd InfixEQSystem implements ExprICG {     EquationSystem = &lt;equations&gt; List(Equation).     List(S) ~ "(" {S} ")".     Equation = &lt;lhs&gt; Variable "=" &lt;rhs&gt; Expression.     Expression : Simple   Compound.     Simple : Variable   Numerical.     Variable = Ident.     Numerical = &lt;v&gt; int.     Compound = &lt;lrand&gt;List(Expression) &lt;op&gt; Op &lt;rrand&gt;                 &gt;List(Expression).     Op : Add.     Add = "+".      use EquationSystem as ESystem,         (-&gt;*,lhs,* to Variable) as DThing,         (-&gt;*,rhs,* to Variable) as UThing,         Variable as Thing. } </pre>
---	---

Figure 3: InfixEQSystem defines a class graph and a mapping of the entities in the class graph to the interface class graph of ExprICG. The driver class Main uses the strategies defined in ExprICG and the IDPrintVisitor.

fact that we can reach more than one variable through the left hand side of the equal sign. With one argument functions the meaning of what is defined and what is its scope has changed and these changes have to be reflected in the Demeter Interface.

It is essential that for this evolution step the interface has to change (Figure 4). With a new interface class graph ParamExprICG we can abstractly reason about semantically checking systems with one argument functions. We also have to strengthen the strategy for the definedIdent traversal specification by adding an extra directive "via Thing". Because definedIdent is a sub-strategy of gdefinedIdent, we need to update gdefinedIdent in the same way. As a final step we update the adaptive methods as shown in Figure 4.

In this evolution step, the Demeter Interface helped by disallowing a naive extension that would violate the intended behavior of the original Demeter Interface. The nature of the evolution required an extension of the interface and that resulted to changes in the driver class and a new concrete class dictionary. It is important to note how the Demeter Interface exposed the erroneous usage of the ExprICG interface for this evolution step and assisted in updating all the dependent components due to the definition of ParamExprICG.

### 2.3 Demeter Interfaces and XPath

Consider the following DTD:

```

<?xml version="1.0"?>
<!ELEMENT BusRoute ((DieselBus|GasBus)*,Town*)>
<!ELEMENT DieselBus (Person*)>
<!ELEMENT GasBus (Person*)>
<!ELEMENT Town (BusStop*)>
<!ELEMENT BusStop (Person*)>
<!ELEMENT Person EMPTY>
<!ATTLIST BusRoute name CDATA #REQUIRED>
<!ATTLIST DieselBus name CDATA #REQUIRED>
<!ATTLIST GasBus name CDATA #REQUIRED>
<!ATTLIST Town name CDATA #REQUIRED>
<!ATTLIST BusStop name CDATA #REQUIRED>
<!ATTLIST Person name CDATA #REQUIRED>

```

and the JavaScript statement containing the XPath expression:

```
var nodes=xmlDoc.selectNodes(" ../BusStop//Person
```

selecting the people that are either waiting at a bus stop or riding in a diesel bus.

Now consider the following DTD:

```

<?xml version="1.0"?>
<!ELEMENT BusRoute ((DieselBus|GasBus)*,Town*)>
<!ELEMENT DieselBus (Person*)>
<!ELEMENT GasBus (Person*)>
<!ELEMENT Town (BusStop*)>
<!ELEMENT BusStop (Person*, CoffeeDrinkers)>
<!ELEMENT CoffeeDrinkers (Person*)>
<!ELEMENT Person EMPTY>
<!ATTLIST BusRoute name CDATA #REQUIRED>
<!ATTLIST DieselBus name CDATA #REQUIRED>
<!ATTLIST GasBus name CDATA #REQUIRED>
<!ATTLIST Town name CDATA #REQUIRED>
<!ATTLIST BusStop name CDATA #REQUIRED>
<!ATTLIST Person name CDATA #REQUIRED>

```

where we added coffee drinkers that are hanging out at a bus stop but who are not waiting for a bus. The JavaScript program still "works" but it gives us the wrong result because no interface exists to protect us. In a Demeter interface we would use:

```
merge(
    from BusRoute via BusStop via -> *,waiting,* to Person
    from BusRoute via DieselBus to Person)

```

to make sure that only the waiting people are selected at bus stops.

For the evolved DTD the interface could be implemented using:

```
var nodes=xmlDoc.selectNodes(" ../BusStop/Person | ../Di

(Notice the single /)
XPath based tools, and there are many of them, all suffer from unprotected XPath expressions using //.
```

### 3. MODULARITY AND DEMETER INTERFACES

The introduction of Demeter Interfaces to AP development assists in designing, maintaining and understanding adaptive programs.

```

di ParamExprICG {
  icg:
    ESystem = List(Definition).
    Definition = <def> DThing "=" <body> Body.
    Body = List(UThing).
    UThing = Thing [ " (" <aparam> UParamName " ) " ].
    DThing = Thing[ " (" <fparam> DParamName " ) " ].
    Thing = Ident.
    UParamName = Ident.
    DParamName = Ident.
    List(S) ~ " ( " S " ) ".

  strategies:
    gdefinedIds = "from ESystem via DThing via Thing to Ident".
    gusedIds = "from ESystem via UThing to Ident".
    defParam = "from Definition via DParamName to Ident".
    definedIdent = "from Definition via DThing via Thing to Ident".

  constraints:
    unique(definedIdent).
    nonempty(gusedIds).
    nonempty(gdefinedIds).
    unique(defParam).

  visitors:
    interface DVisitor {
      public void before (Thing t);
      public LinkedList return ();
    }

    interface PVisitor extends DVisitor {
      public void before (Definition t);
    }

  adaptive methods:
    ESystem {
      public LinkedList printDefined(PVisitor) gdefinedIds;
      public LinkedList printUsed(PVisitor) gusedIds;
      public boolean checkBindngs(LinkedList l1, LinkedList l2){
        // checks appropriate usage of variables
      }
    }
}

```

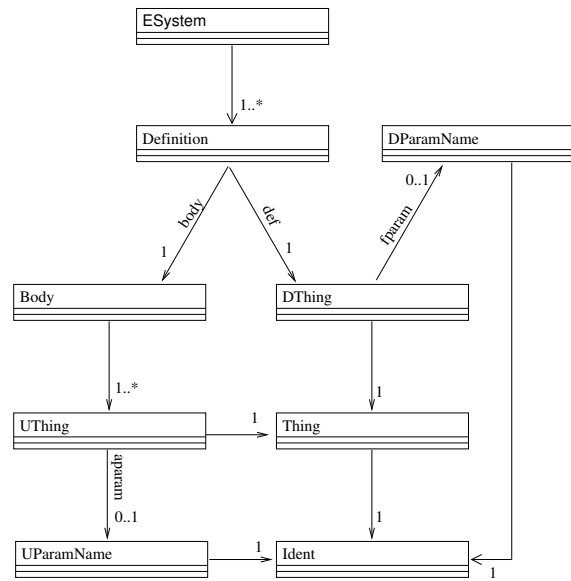


Figure 4: The evolved Demeter Interface (left) and the UML representation of the extended Demeter Interface class graph (right).

Eventhough we have just began to explore the ideas behind DIs their usage has revealed several benefits.

During the design process of adaptive programs developers would first design a minimal class dictionary. Then iteratively both adaptive code and the class dictionary itself are developed with repeated testing to verify the behavior of adaptive code. Modifications to both the class dictionary as well as the adaptive code (the traversal specification and/or the visitor attached) were necessary. As programs become large in size distinguishing which parts of the class dictionary are involved in the different adaptive methods becomes difficult. Further more, modifications to class names in the class dictionary cause changes to traversal strategies and/or visitor methods due to the lack of abstraction over class names. For example, in the developement of CONA [7, 8] a Design by Contract (DbC) extension to Java and AspectJ, the class dictionary is the whole Java and AspectJ language. Understanding the dependencies between adpative code and the class dictionary becomes a laborious and error prone process.

DI provide solutions to both of these problems. The icg pro-

vides an abstraction of the concrete class graph while the adaptive methods, traversal strategies and visitor interfaces localize all the information necessary for understanding the dependencies between adaptive code and the rest of the program. The mapping mechanism removes the tight dependence on naming conventions by providing an automatic renaming mechanism. The usage of DI allows for more modular AP designs.

To support our claim of modularity for DIs we borrow the definition for modular implementations as proposed by Kiczales and Mezini [?]

- it is textually local,
- there is a well-defined interface that describes how it interacts with the rest of the system,
- the interface is an abstraction of the implementation, in that it is possible to make material changes to the implementation without violating the interface.
- an automatic mechanism enforces that every module satisfies its own interface and respects the interface of all other

```

class Main {
  public static void main(String[] args){
    Visitor v = new IDPrintVisitor();
    GVisitor v1 = new GVisitor();
    ParamEquations pe = ParamEquations.parse(System.in);
    boolean codeOk = pe.checkBindings(pe.printDefined(v),
    pe.printUsed(v1));
    if (!codeOk) i
      System.out.println(" Variables used before they where
        defined");
  }
}

```

```

class IDPrintVisitor extends DVisitor {
  public void before(Ident id) {
    System.out.println(id);
  }
}
class GVisitor extends PVisitor {
  public void before(Definition d) {
    Ident dparam = (Ident) ParamEQSystem.fetch(d,
    defParam);
    System.out.println(" defined parameter: " +
    dparam);
  }
  public void before(Ident id) {
    System.out.println(id);
  }
}

```

Figure 5: Changes to the interface affect Main. The definition of GVisitor is used to check for the local parameter names in parametric equations.

modules, and

- the module can be automatically composed – by a compiler, loader, linker etc. – via various configurations with other modules to produce a complete system.

DIs are textually local with traversal strategies and visitors specifying exactly how adaptive methods interact with the rest of the system. DIs are an abstraction of the implementation both of the class dictionary, through the *icg*, but also through the visitor interfaces that are part of the DI's definition. The extensions made to the DemeterJ system provide automatic mechanisms that both check that modules satisfy their own interfaces as well as the interfaces of other modules. Composition of DI automatically managed by the DemeterJ system and configuration of the composition can be controlled via the *implements* keyword and the mapping specification.

#### 4. RELATED WORK

Ovlinger and Wand [?] propose a domain specific language as a means to specify recursive traversals of object structures used with the visitor pattern [?]. The domain specific language further allows for intermediate results from subtraversals allowing for a more functional style visitor definitions. The explicit full definition of the recursive data structure provides an interface between visitors and the underlying data structure. This approach enforces that each object in a traversal is explicitly defined allowing no room for adaptiveness.

Modularity issues in AOSD [2, 1, 9] have received great attention recently. Kiczales and Mezini [2] advocate that in the presence of aspects, a module's interface has to further include pointcuts from aspects that apply to the module in question. These augmented interface definitions, named *aspect-aware interfaces*, can only be determined after the complete configuration of the system's components is known. Aspect-aware interfaces do not provide any extra information hiding capabilities to the base program's modules.

Open Modules [1], extend the traditional notion of a software module to include in its interface pointcut specifications. In this way a module can export, and as such make publicly available, pointcuts within its implementation. This approach gives a balanced control between module and aspect developers in terms of information hiding thus allowing for separate (parallel) evolution of aspects and modules on the agreed upon interface. The interface

of a crosscutting concern can affect multiple modules at different join points on each one. Thus an aspect's interface is sprinkled along module interfaces and not localized making it harder (if not impossible at times) for aspect developers to develop their aspects.

Kevin Sullivan and Bill Griswold and their students [9] advocate an XPI (crosscutting program interface) as a means to achieve separate development and explicit dependencies between implementations of crosscutting concerns and base code. XPIs at present are a design artifact where the agreed upon interface between aspects and base code is explicitly stated using English. Although XPIs assist both during design and development, there is no mechanical checking available to verify the implementation against the agreed upon interface.

Kiczales and Mezini in [3] discuss the benefits of using different programming language mechanisms (procedures, annotations, advice and pointcuts) used to provide separation of concerns at the code level. The resulting guidelines from their analysis sketch the situations where each mechanism will be most effective. The inherent modularity issues associated with each technology are not addressed.

#### 5. CONCLUSION

We introduce Demeter Interfaces as an extension to DemeterJ. Demeter interfaces encapsulate all the information and dependencies between the adaptive code and the underlying data structure. Through the definition of an interface class graph and a set of graph constraints Demeter Interfaces impose restrictions on any concrete data structure to which adaptive code will be attached. These restrictions are enforced at compile time disallowing modifications to the underlying data structure that would otherwise provide incorrect program results. With Demeter interfaces in place we have shown how modularity as well as understandability of adaptive programs increases dramatically leading to better program design and promotes parallel development.

#### 6. FUTURE WORK

As immediate future work we would like to complete our implementation of Demeter Interfaces and study extensions to our ideas that would better facilitate AP programming. The set of available predicates that can be used as constraints is still under investigation. Eventhough with the available set of predicates we could capture all of the constraints that we needed, a bigger set of AP programs needs to be investigated. Some recommendations have been made to

include set operators (like subset, proper subset) that can be used on the set of classes that are captured by traversals. Also being able to compose these predicates through logical operators into boolean expressions is a plan for the near future.

Another extension would be the provision of contracts on adaptive methods. Collaborations of adaptive methods exchange data and certain assumptions are being made that are not explicitly captured by the current system. The addition of pre- and post-conditions would assist in both defining and validating these assumptions. The composition of Demeter Interfaces to provide new Demeter Interfaces is also an interesting future research direction. Also the interactions between DI in an AP program still remains an open issue.

A long term goal would be the extension of ideas behind Demeter Interfaces to more general purpose AOP languages. Specifically exposing a graph based model for the control flow of programs upon which *control interfaces* of their expected behavior can be defined. Each program module, *e.g.*, Java class or Java package, will be accompanied by a control interface that will specify an abstraction of the call graph within that module along with constraints. Aspects that are to be attached to this module will have to do so based on the module's interface. A composition of two modules will generate a composition of their corresponding control interfaces. Aspect developers and class developers will first agree on the control interface of their modules and the assumptions that it must uphold. The implementation of the module and the aspects around it will then be checked against the agreed upon control interface. Similar ideas for interfaces between aspects and base code have been recently discussed [1, 9, 3].

## 7. REFERENCES

- [1] Jonathan Aldrich. Open Modules: modular reasoning about advice. In *European Conference on Object-Oriented Programming*, 2005.
- [2] Matthias Felleisen, Robert Bruce Findler, Matthew Flatt, and Shriram Krishnamurthi. *How to Design Programs*. MIT Press, 2001.
- [3] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley, 1995.
- [4] Gregor Kiczales and Mira Mezini. Aspect-oriented programming and modular reasoning. In *ICSE '05: Proceedings of the 27th International Conference on Software Engineering*, pages 49–58, New York, NY, USA, 2005. ACM Press.
- [5] Gregor Kiczales and Mira Mezini. Separation of concerns with procedures, annotations, advice and pointcuts. In *European Conference on Object-Oriented Programming*, 2005.
- [6] Karl J. Lieberherr. *Adaptive Object-Oriented Software: The Demeter Method with Propagation Patterns*. PWS Publishing Company, Boston, 1996. 616 pages, web book at [www.ccs.neu.edu/research/demeter](http://www.ccs.neu.edu/research/demeter).
- [7] Karl J. Lieberherr. *Adaptive Object-Oriented Software: The Demeter Method with Propagation Patterns*. PWS Publishing Company, Boston, 1996. 616 pages, ISBN 0-534-94602-X. 
- [8] Doug Orleans and Karl J. Lieberherr. DemeterJ. Technical report, Northeastern University, 1996-2001. <http://www.ccs.neu.edu/research/demeter/DemeterJava/>.
- [9] Johan Ovlinger and Mitchell Wand. A Language for Specifying Traversals of Object Structures. Technical Report NU-CCS-98-??, College of Computer Science, Northeastern University, Boston, MA, November 1998.
- [10] Therapon Skotiniotis and David Lorenz. Conaj: Generating contracts as aspects. Technical Report NU-CCIS-04-03, College of Computer and Information Science, Northeastern University, 2004.
- [11] Therapon Skotiniotis and David H. Lorenz. Cona: aspects for contracts and contracts for aspects. In *OOPSLA '04: Companion to the 19th annual ACM SIGPLAN conference on Object-oriented programming systems, languages, and applications*, pages 196–197, New York, NY, USA, 2004. ACM Press.
- [12] Kevin Sullivan, William G. Griswold, Yuanyuan Song, Yuanfang Cai, Macneil Shonle, Nisit Tewan, and Hridesh Rajan. On the criteria to be used in decomposing systems into aspects. In *European Software Engineering Conference and International Symposium on the Foundations of Software Engineering*, 2005.