
ASPECT-ORIENTED PROGRAMMING

GREGOR KICZALES, JOHN LAMPING, ANURAG MENDHEKAR, CHRIS MAEDA,
CRISTINA LOPES, JEAN-MARC LOINGTIER AND JOHN IRWIN

©Copyright 1997, XEROX Corporation. All rights reserved.

XEROX PALO ALTO RESEARCH CENTER*

We have found many programming problems for which neither procedural nor object-oriented programming techniques are sufficient to clearly capture some of the important design decisions the program must implement. This causes the implementation of those design decisions to be scattered throughout the code, resulting in “tangled” code that is excessively difficult to develop and maintain. We present an analysis of why it is that such design decisions have been so difficult to clearly capture in actual code. We call the issues these decisions address ASPECTS, and say that the reason they have been hard to capture is that they CROSS-CUT the system’s basic functionality. We present the basis for a new programming technique, called aspect-oriented programming, that makes it possible to clearly express programs involving such aspects, including appropriate isolation, composition and reuse of the aspect code. The discussion is rooted in systems we have built using aspect-oriented programming.

1. INTRODUCTION

Object-oriented programming has been presented as a technology that can fundamentally aid software engineering, because the underlying object model provides a better fit with real domain problems. But we have found many programming problems where OOP techniques are not sufficient to clearly capture all the important design decisions the program must implement. Instead, it seems that there are some programming problems that fit neither the OOP approach nor the procedural approach it replaces.

This paper reports on our work developing programming techniques that make it possible to clearly express those programs that OOP fails to support. We present an analysis of why some design decisions have been so difficult to cleanly capture in actual code. We call the issues these decisions address ASPECTS, and say that the reason they have been hard to capture is that they CROSS-CUT the system’s basic functionality. We present the basis for a new programming technique, called aspect-oriented pro-

* 3333 Coyote Hill Road, Palo Alto, CA 94304, USA. gregor@parc.xerox.com

gramming (AOP), that makes it possible to clearly express programs involving such aspects, including appropriate isolation, composition and reuse of the aspect code.

We think of the current state of AOP research as analogous to that of OOP 20 years ago. The basic concepts are beginning to take form, and an expanding group of researchers are using them in their work [1, 2, 3]. Furthermore, while AOP qua AOP is a new idea, there are existing systems that have AOP-like properties. The contribution of this paper is an analysis of the problems AOP is intended to solve, as well as an initial set of terms and concepts that support explicit AOP-based system design.

The paper presents AOP in an example-driven way—the generalizations and definitions are all derived from examples, rather than presented in advance. Section 3 use a medium-scale example to present the aspect-tangling problem AOP solves; the section culminates with a definition of the term aspect. Section 4 presents several more small examples of aspects. Sections 5 and 6 each provide an example of a complete AOP system. The remaining sections present future work, related work and conclusions.

2. BACKGROUND ASSUMPTIONS

This section outlines important assumptions about the relationship between programming languages and software design processes that underlie the rest of the paper.

Software design processes and programming languages exist in a mutually supporting relationship. Design processes break a system down into smaller and smaller units. Programming languages provide mechanisms that allow the programmer to define abstractions of sub-units of their systems, and then compose those abstractions in different ways to produce the overall system. A design process and a programming language work well together when the programming language provides abstraction and composition mechanisms that cleanly support the units the design process breaks the system into.

From this perspective, many existing programming languages, including object-oriented languages, procedural languages and functional languages, can be seen as having a common root in that their key abstraction and composition mechanisms are all rooted in some form of generalized procedure. For the purpose of this paper we will refer to these as generalized-procedure (GP) languages. (This is not to say that we are ignorant of the many important advantages of OOP languages! It is only to say that for the purposes of the discussion in this paper, it is simpler to focus on what is common across all GP languages.)

The design methods that have evolved to work with GP languages tend to break systems down into units of behavior or function. This style has been called functional decomposition [4,5,6].¹ The exact nature of the decomposition differs between the language paradigms of course, but each unit is encapsulated in a procedure/function/object, and in each case, it feels comfortable to talk about what is encapsulated as a *functional* unit of the overall system. This last point may be so familiar that it feels somewhat redundant. But it is important that we give it explicit attention now, because in the course of this paper will be considering units of system decomposition that are not functional.

¹ In some parts of the computing community this term connotes the use of functional programming languages (i.e. side-effect free functions), but in this paper we do not use the term in that sense.

3. WHAT ARE ASPECTS?

To better understand the origins of tangling problems, and how AOP works to solve them, this section is organized around a detailed example program. This program represents a simplified version of a real application we have been working with [7,8]—the tangling we discuss is real (but simplified); the AOP solution presented in Section 5 is real (but simplified).

Consider the implementation of a black-and-white image processing system, in which the desired domain model is one of images passing through a series of filters to produce some desired output. Assume that important goals for the system are that it be easy to develop and maintain, and that it make efficient use of memory. The former because of the need to quickly develop bug-free enhancements to the system. The latter because the images are large, so that in order for the system to be efficient, it must minimize both memory references and overall storage requirements.

3.1. BASIC FUNCTIONALITY

Achieving the first goal is relatively easy. Good old-fashioned procedural programming can be used to implement the system, concisely, and in good alignment with the domain model. In such an approach the filters can be defined as procedures that take several input images and produce a single output image. A set of primitive procedures would implement the basic filters, and higher level filters would be defined in terms of the primitive ones. For example, a primitive `or!` filter, which takes two images and returns their pixelwise logical or, might be implemented as:²

```
(defun or! (a b)
  (let ((result (new-image)))
    (loop for i from 1 to width do
      (loop for j from 1 to height do
        (set-pixel result i j
          (or (get-pixel a i j)
              (get-pixel b i j))))))
  result))
```

the operation to perform on the pixels ——— (or (get-pixel a i j) (get-pixel b i j)))) ——— storing pixels in the result image

loop over all the pixels in the input images ——— (loop for i from 1 to width do (loop for j from 1 to height do

Starting from `or!` and other primitive filters, the programmer could work up to the definition of a filter that selects just those black pixels on a horizontal edge, returning a new image consisting of just those boundary pixels.

² We have chosen Common Lisp syntax for this presentation, but this could be written fairly easily in any other Algol-like language.

functionality

pixelwise logical operations
shift image up, down
difference of two images
pixels at top edge of a region
pixels at bottom edge of a region
horizontal edge pixels

implementation

written using loop primitive as above

written using loop primitive; slightly different loop structure

```
(defun remove! (a b) (and! a (not! b)))  
(defun top-edge! (a) (remove! a (down! a)))  
(defun bottom-edge! (a) (remove! a (up! a)))  
(defun horizontal-edge! (a)  
  (or! (top-edge! a) (bottom-edge! a)))
```

Note that only the primitive filters deal explicitly with looping over the pixels in the images. The higher level filters, such as `horizontal-edge!`, are expressed cleanly in terms of primitive ones. The resulting code is easy to read, reason about, debug and extend—in short, it meets the first goal.

3.2. OPTIMIZING MEMORY USAGE

But this code doesn't address the second goal of optimizing memory usage. When each procedure is called, it loops over a number of input images and produces a new output image. Output images are created frequently, often existing only briefly before they are consumed by some other loop. This results in excessively frequent memory references and storage allocation, which in turn leads to cache misses, page faults, and terrible performance.

The familiar solution to the problem is to take a more global perspective of the program, map out what intermediate results end up being inputs to what other filters, and then code up a version of the program that fuses loops appropriately to implement the original functionality while creating as few intermediate images as possible. The revised code for `horizontal-edge!` would look something like:

```
(defun horizontal-edge! (a)  
  (let ((result (new-image))  
        (a-up (up! a))  
        (a-down (down! a)))  
    (loop for i from 1 to width do  
      (loop for j from 1 to height do  
        (set-pixel result i j  
          (or (and (get-pixel a i j)  
                  (not (get-pixel a-up i j)))  
              (and (get-pixel a i j)  
                  (not (get-pixel a-down i j)))))))  
    result))
```

only three result images are created

one loop structure shared by many component filters

operations from many component filters

Compared to the original, this code is all tangled up. It incorporates all the different filters that `horizontal-edge!` is defined in terms of, and fuses many, but not all, of their loops together.

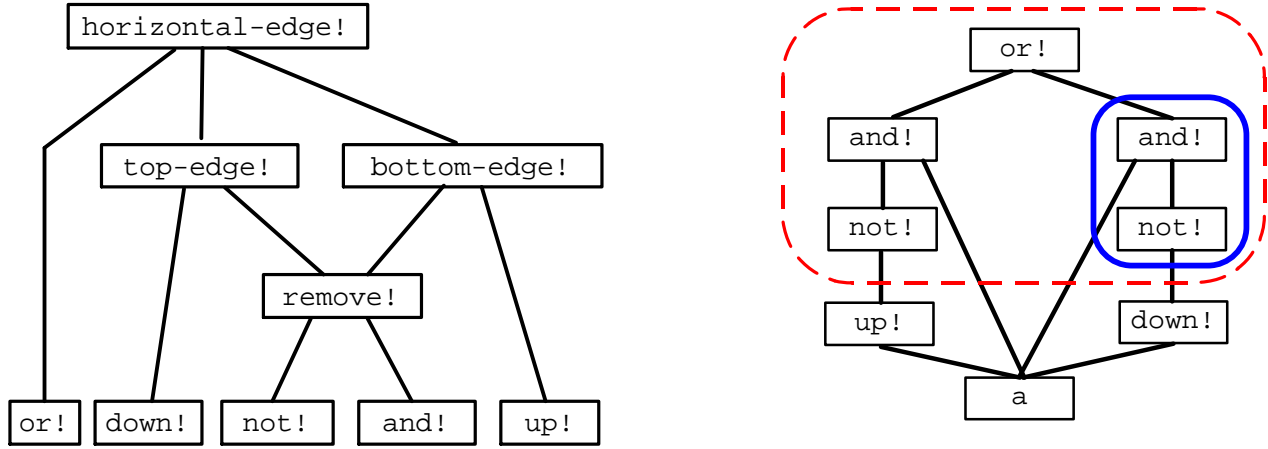


Figure 1: Two different diagrams of the un-optimized `horizontal-edge!` filter. On the left is the hierarchical procedure calling structure, which aligns so directly with the domain model. On the right is a data flow diagram, in which the boxes are the primitive filters and the edges are the data flows between them at runtime. The box labeled `a` at the bottom is the input image. The red dashed oval indicates the primitives that are fused together in the optimized version. To help understand the relation between the diagrams, the blue solid oval indicates the functional boundary of one invocation of `remove!` during the execution.

(The loops for `up!` and `down!` are not fused because those operations have a different looping structure.)³ In short, revising the code to make more efficient use of memory has destroyed its original clear functional structure.

Of course, this is a very simple example, and it is not so difficult to deal with such a small amount of tangled code. But in real programs the complexity due to such tangling quickly expands to become a major obstacle to ease of code development and maintenance. The real system this example was drawn from is an important sub-component of an optical character recognition system. For that real system, coding the basic functionality in the simple style requires only 768 lines of code; but the optimized version, which does the fusion optimization as well as memoization of intermediate results, compile-time memory allocation and specialized intermediate datastructures, requires 35213 lines. The tangled code is extremely difficult to maintain, since small changes to the functionality require mentally untangling and then re-tangling it.

3.3. CROSS-CUTTING

Returning to the example code, Figure 1 provides a different basis for understanding the tangling in it. On the left there is the hierarchical structure of the filtering functionality. On the right there is a data flow diagram for the original, un-optimized version of `horizontal-edge!`. In this diagram, the boxes and lines show the primitive filters and data flow between them. The dashed oval shows the boundary of what is fused into a single loop in the optimized version of `horizontal-edge!`.

³ Fusing these other loops as well is not so difficult. We chose not to show that code because it is so tangled that it is distractingly difficult to understand.

Notice that the fusion oval does not incorporate all of horizontal-edge! In fact, it doesn't align with any of the hierarchical units on the left. While the two properties being implemented—the functionality and the loop fusion—both originate in the same primitive filters, they must compose differently as filters are composed. The functionality composes hierarchically in the traditional way. But the loop fusion composes by fusing the loops of those primitive filters that have the same loop structure and that end up in producer/consumer relationships at runtime. This fusion relationship cuts across the functionality hierarchy so fundamentally that the only way to make it appear local is to look at the dataflow graph.

This cross-cutting phenomena is directly responsible for the tangling in the code. The single composition mechanism the language provides us—procedure calling—is very well suited to building up the un-optimized functional units. But it can't help us compose the functional units and the loop fusion simultaneously, because they follow such different composition rules and yet must co-compose. Unfortunately, this breakdown forces us to combine the properties entirely by hand—that's what happening in the tangled code above.

In general, whenever two properties being programmed must compose differently yet be coordinated, we say that they CROSS-CUT each other. Because GP languages provide only one composition mechanism, the programmer must do the co-composition manually, leading to complexity and tangling in the code.

We can now define two important terms more precisely:

With respect to a system and its implementation using a GP-based language, a property that must be implemented is:

A COMPONENT, if it can be cleanly encapsulated in a generalized procedure (i.e. object, method, procedure, API). By cleanly, we mean well-localized, and easily accessed and composed as necessary. Components tend to be units of the system's functional decomposition, such as image filters, bank accounts and GUI widgets.

An ASPECT, if it can not be cleanly encapsulated in a generalized procedure. Aspects tend not to be units of the system's functional decomposition, but rather to be properties that affect the performance or semantics of the components in systemic ways. Examples of aspects include memory access patterns and synchronization of concurrent objects. (Section 4 provides more examples of aspects.)

Using these terms it is now possible to clearly state the goal of AOP: To support the programmer in cleanly separating components and aspects from each other,⁴ by providing mechanisms that make it possible to abstract and compose them to produce the overall system. This is in contrast to GP-based programming, which supports programmers in separating only components from each other by providing mechanisms that make it possible to abstract and compose them to produce the overall system.

⁴ Components from each other, aspects from each other, and components from aspects.

4. OTHER EXAMPLES OF HOW ASPECTS CROSS-CUT COMPONENTS

Before going on to the presentation of AOP, and how it solves the problem of aspect tangling in code, this section briefly presents several more examples of aspects and components. For each example in the table below we list an application, a kind of GP language that would do a good job of capturing the component structure of the application, a likely component structure for the application if programmed using that kind of language, and the aspects that would cross-cut that component structure.

application	GP language	components	aspects
image processing	procedural programming	filters	loop fusion result sharing compile-time memory allocation
digital library	object-oriented programming	catalogs, printers, services	communication strategy synchronization constraints failure handling
matrix algorithms	procedural programming	linear algebra operations	matrix representation permutation floating point error

Some aspects are so common that they can easily be thought about without reference to any particular domain. One of the best examples is error and failure handling. We are all familiar with the phenomenon that adding good support for failure handling to a simple system prototype ends up requiring many little additions and changes throughout the system. This is because the different dynamic contexts that can lead to a failure, or that bear upon how a failure should be handled, cross-cut the functionality of systems.

Many performance-related issues are aspects, because performance optimizations often exploit information about the execution context that spans components.

In this section we return to the image processing example, and use it to sketch an AOP-based implementation of that application. The presentation is based on a system we have developed, but is simplified somewhat. The complete system is discussed in [7]. The goal of this section is to quickly get the complete structure of an AOP-based implementation on the table, not to fully explain that structure. Section 6 will provide that explanation.

The structure of the AOP-based implementation of an application is analogous to the structure of a GP-based implementation of an application. Whereas a GP-based implementation of an application consists of: (i) a language, (ii) a compiler (or interpreter) for that language, and (iii) a program written in the language that implements the application; the AOP-based implementation of an application consists of: (i.a) a COMPONENT LANGUAGE with which to program the components, (i.b) one or more ASPECT LANGUAGES with which to program the aspects, (ii) an ASPECT WEAVER™ for the combined languages, (iii.a) A COMPONENT PROGRAM, that implements the components using the component language, and (iii.b) one or more ASPECT PROGRAMS that implement the aspects using the aspect languages. Just as with GP-based languages, AOP languages and weavers can be designed so that weaving work is delayed until runtime (RT weaving), or done at compile-time (CT weaving).

5.1. THE COMPONENT LANGUAGE & PROGRAM

In the current example we use one component language and one aspect language. The component language is similar to the procedural language used above, with only minor changes. First, filters are no longer explicitly procedures. Second, the primitive loops are written in a way that makes their loop structure as explicit as possible. Using the new component language the `or!` filter is written as follows:

```
(define-filter or! (a a)
  (pixelwise (a b) (aa bb) (or aa bb)))
```

The `pixelwise` construct is an iterator, which in this case walks through images `a` and `b` in lock-step, binding `aa` and `bb` to the pixel values, and returning a image comprised of the results. Four similar constructs provide the different cases of aggregation, distribution, shifting and combining of pixel values that are needed in this system. Introducing these high-level looping constructs is a critical change that will enable the aspect languages to be able to detect, analyze and fuse loops much more easily.

5.2. THE ASPECT LANGUAGE & PROGRAM


The design of the aspect language used for this application is based on the observation that the dataflow graph in Figure 1 makes it easy to understand the loop fusion required. The aspect language is a simple procedural language that provides simple operations on nodes in the dataflow graph. The aspect program can then straightforwardly look for loops that should be fused, and carry out the fusion required. The following code fragment is the core of the code that handles the fusion shown above. It checks whether two nodes, that are connected by a data flow edge, both have a `pixelwise` loop structure, and if so it fuses them into a single loop that also has a `pixelwise` structure, and that has the appropriate merging of the inputs, loop variables and body of the two original loops.


```
(cond ((and (eq (loop-shape node) 'pointwise)
            (eq (loop-shape input) 'pointwise))
      (fuse loop input 'pointwise :inputs (splice ...)
                                :loop-vars (splice ...)
                                :body (subst ...))))
```

Describing the composition rules and fusion structure for the five kinds of loops in the real system requires about a dozen similar clauses about when and how to fuse.⁵ This is part of why this system could not be handled by counting on an optimizing compiler to do the appropriate fusion—the program analysis and understanding involved is so significant that compilers cannot be counted upon to handle this reliably. (Although many compilers might be able to optimize this particular simple example.) Another complication is the other aspects the real system handles, including sharing of intermediate results and keeping total runtime memory allocation to a fixed limit.

5.3. WEAVING

The aspect weaver accepts the component and aspect programs as input, and emits a C program as output. This work proceeds in three distinct phases, as illustrated in the following figure and explained below it.



In phase 1 the weaver uses unfolding as a technique for generating a data flow graph from the component program. In this graph, the nodes represent primitive filters, and the edges represent an image flowing from one primitive filter to another. Each node contains a single loop construct. So, for example, the node labeled **A** contains the following loop construct, where the `#<...>` refer to the edges coming into the node:

```
(pointwise (#<edge1> #<edge2>) (i1 i2) (or i1 i2))
```

In phase 2 the aspect program is run, with access to the graph. The effect of this is to edit the graph by collapsing nodes together and adjusting their bodies accordingly. The result is a graph in which some of the loop structures have more primitive pixel operations in them than before phase 2. For ex-

⁵ Different development versions of the real system have handled different amounts of the total possible fusion. As described in [7], the most recent version does not handle this additional fusion, we are adding that in current development.

ample, the node labeled **B**, which corresponds to the fusion of 5 loops from the original graph, has the following as its body:

```
(pointwise (#<edge1> #<edge2> #<edge3>) (i1 i2 i3)
  (or (and (not i1) i2) (and (not i3) i2))))
```

Finally, in phase 3, a simple code generator walks over the fused graph, generating one C function for each loop node, and generating a main function that calls the loop functions in the appropriate order, passing them the appropriate results from prior loops. The code generation is simple because each node contains a single loop construct with a body composed entirely of primitive operations on pixels.

A crucial feature of this system is that the weaver is not a “smart” compiler, which can be so difficult to design and build. By using AOP, we have arranged for all the significant implementation strategy decisions—all the actual smarts—to be provided by the programmer, using the appropriate aspect languages. The weaver’s job is integration, rather than inspiration.⁶

5.4. RESULTS

The real system is somewhat more complex of course. For one thing, there are two additional aspect programs, one of which handles sharing of common sub-computations, and one of which ensures that the minimum possible number of images are allocated at any one time. All three of the aspect programs are written in the same aspect language.

In this example, the AOP approach has met the design goals—the target application code is easy to reason about, develop and maintain, while at the same time being highly efficient. It is easy for the programmer to understand the components and how they compose. It is easy for the programmer to understand the aspects, and how they compose. It is easy for the programmer to understand the effect of the aspect programs on the total output code. Changes in either the filter components or the fusion aspect are easily propagated to the whole system by simply re-weaving. What isn’t easy is for the programmer to generate the details of the output code. The power of the AOP approach is that the weaver handles these details, instead of the programmer having to do the tangling manually.

Our AOP based re-implementation of the application is 1039 lines of code, including the component program and all three aspect programs. The aspect weaver itself, including a reusable code generation component is 3520 lines (the true kernel of the weaver is 1959 lines). Its performance is comparable to a 35213 lines manually tangled version produced by a separate project (the time efficiency is comparable and the space efficiency is better).

As with many other software engineering projects, it is extremely difficult to quantify the benefits of using AOP without a large experimental study, involving multiple programmers using both AOP and traditional techniques to develop and maintain different applications [13, 14, 15]. Such a study has been

⁶ While asking the programmer to explicitly address implementation aspects sounds like it might be a step backwards, our experience with work on open implementation suggests that in fact it isn’t [9, 10, 11, 12]. While the programmer is addressing implementation in the memory aspect, proper use of AOP means that they are expressing implementation strategy at an appropriately abstract level, through an appropriate aspect language, with appropriate locality. They are not addressing implementation details, and they are not working directly with the tangled implementation.

beyond the scope of our work to date, although we hope to do one in the future. In the meantime, we have developed one initial measure of the degree to which applying AOP techniques can simplify an application. This measure compares a GP-based implementation of an application to an AOP-based implementation of the same application. It measures the degree to which the aspects are more concisely coded in the AOP-based implementation than in the non AOP-based implementation. The general equation for this measure, as well as the numbers for this particular application are as follows:

$$\text{reduction in bloat due to tangling} = \frac{\text{tangled code size} - \text{component program size}}{\text{sum of aspect program sizes}} = \frac{35213 - 756}{352} = 98$$

In this measure, any number greater than 1 indicates a positive outcome of applying AOP. Note that as a single number, this measure has somewhat limited utility—it is useful in this case because on the other important grounds, namely performance, the AOP-based implementation of the application is comparable to the non-AOP based implementation. Section 7 presents some of the requirements we have identified for quantitative measures of AOP utility.

6. SECOND EXAMPLE OF AOP

This section uses a second example of an AOP-based system to elaborate on component language design, aspect language design and weaving. Once again, the example is a simplified version of a real system we are developing, which is described in [16]. The example comes from the document processing domain where we wanted to implement a distributed digital library that stores documents in many forms and provides a wide range of operations on those documents. The AOP system developed in this section is more general-purpose in nature than the highly domain-specific example in the previous section.

The functionality of this system is well captured using an object-oriented model. In such an approach the objects are things like documents, catalogs, different printable forms for the documents (pdf, ps, rip...), printers, servers etc. There are several aspects of concern, including:

- Communication, by which we mean controlling the amount of network bandwidth the application uses by being careful about which objects and sub-objects get copied in remote method calls. For example, we want to be sure that when a book object is included in a remote method invocation, the different printed representations of the book aren't sent across the wire unless they will be needed by the receiving method.
- Coordination constraints, by which we mean the synchronization rules required to ensure that the component program behaves correctly in the face of multiple threads of control.
- Failure handling, by which we mean handling the many different forms of failure that can arise in a distributed system in an appropriately context-sensitive way.

For now, we will continue with just the communication aspect. Handling both communication and coordination using AOP is discussed in [16]. Failure handling using AOP is a future research goal.

6.1. THE COMPONENT LANGUAGE & PROGRAM

Designing an AOP system involves understanding what must go into the component language, what must go into the aspect languages, and what must be shared among the languages. The component language must allow the programmer to write component programs that implement the system's functionality, while at the same time ensuring that those programs don't pre-empt anything the aspect programs need to control. The aspect languages must support implementation of the desired aspects, in a natural and concise way. The component and aspect languages will have different abstraction and composition mechanisms, but they must also have some common terms, these are what makes it possible for the weaver to co-compose the different kinds of programs.

In the image processing system, replacing low-level loops with the higher-level looping primitives is an example of a change that ensures that component programs don't pre-empt aspect programs. This change makes it easier for the aspect programs to detect and implement opportunities for loop fusion.

In this example, component programs must implement elements such as books, catalogs, and printers. In order to allow the communication aspect program to handle communication, component programs must avoid doing so. In this case Java™ serves quite well as the component language. It provides an object model that implements the appropriate components, and avoids addressing the communication aspect.⁷ So, using Java as our component language, the definition of two simple classes, books and catalogs of books, look like:

⁷ [16] explains that in order to support the coordination aspect language, some lower-level synchronization features must be removed from Java before it can be used as the component language. These are the keyword `synchronized`, and the methods `wait`, `notify` and `notifyAll`.

<pre> public class Book { String title, author; int isbn; OCR ocr; PDF pdf; Postscript ps; RIP rip; public String get_title() { return title; } public String get_author() { return author; } public int get_isbn() { return isbn; } } </pre>	<pre> public class Catalog { private Book books[]; private int nbooks = 0; public Catalog (int dbsize) { books = new Book[dbsize]; } public void register (Book b) { books[nbooks++] = b; } public void unregister (Book b) { ... } public Book lookup (String str) { ... } } </pre>
--	--

6.2. THE ASPECT LANGUAGE & PROGRAM

Communication aspect programs would like to be able to control the amount of copying of arguments that takes place when there is a remote method invocation. To do this, the aspect language must effectively allow them to step into the implementation of method invocation, to detect whether it is local or remote, and to implement the appropriate amount of copying in each case.

One way to do this is to provide runtime reflective access to method invocation. As has been shown in [17,18,19,20] such reflective access can be used to control the communication aspect of a distributed object system. But this kind of reflective access is so powerful that it can be dangerous or difficult to use. So in this case we have chosen to provide a higher-level aspect language, that is more tailored to the specific aspect of controlling copying in remote method invocations. This aspect language can be seen as being built on top of the lower level one, in the same way as it is common to build higher-level GP-based languages on top of lower-level ones.

The communication aspect language we have designed allows the programmer to explicitly describe how much of an object should be copied when it is passed as an argument in a remote method invocation. Using that language, the following fragment of the communication aspect program says that when books are registered with a catalog, only the ISBN number and title should be copied in the call, when they are de-registered, only the ISBN number is copied. The rest of the book, including large sub-objects such as the printable representations, is not copied unless it is needed at some later time.

```

remote Catalog {
    void register (Book: copy isbn, copy title);
    void unregister (Book: copy isbn);
    Book: copy isbn lookup(String);
}

```

6.3. ASPECT WEAVER

Aspect weavers must process the component and aspect languages, co-composing them properly to produce the desired total system operation. Essential to the function of the aspect weaver is the concept of JOIN POINTS, which are those elements of the component language semantics that the aspect programs coordinate with.

In the image processing example, the join points are the nodes and edges in the data flow graph of the component program. In this distributed objects example, the join points are the runtime method invocations in the component program. These two examples serve to illustrate an important point about join points—they are not necessarily explicit constructs in the component language. Rather, like nodes in the dataflow graph and runtime method invocations they are clear, but perhaps implicit in the source code, elements of the component program's semantics.

Aspect weavers work by generating a JOIN POINT REPRESENTATION of the component program, and then executing (or compiling) the aspect programs with respect to it. In the digital library example, the join-point representation includes information about dynamic method invocations such as the concrete classes of the arguments and their location. The communication aspect language is implemented as an interpreter, called at each method invocation, which uses the join point information to appropriately marshal the arguments.⁸ Thus the higher-level aspect language we have designed is implemented on top of a lower level one, as often happens in GP languages.

In the image processing application, the join point representation is quite simple. It is just the data flow graph, operations to access the body of nodes, and operations to edit the graph.

7. OPEN ISSUES

As an *explicit* approach to programming, AOP is a young idea. Our work to date has been primarily focused on designing and implementing aspect-oriented programming languages, and using those languages to develop prototype applications. This programming-centric initial focus has been natural, and it parallels the early development of OOP. But there is a great deal of work still to be done to assess the overall utility of AOP, to better understand its relation to existing ideas, and to further develop it so that it can be useful for a wide range of users.

One important goal is quantitative assessment of the utility of AOP. How much does it help in the development of real-world applications? How much does it help with maintenance? Can we develop measures of which applications it will be more or less useful for? This is a difficult problem, for all the same reasons that quantitative assessment of the value of OOP has been difficult, but we believe that it is important to begin work on this, given that it will take time to get solid results.

We also believe it is important to begin a systematic study to find existing systems that have AOP-like elements in their design. We see this as a way to quickly accelerate development of the AOP ideas,

⁸ In the actual system the weaver is split into compile-time and runtime parts, so that much of the work is done at compile time.

by providing a way to get rough empirical evidence without having to build large new systems from the ground up.

Another important area for exploration is the space of different kinds of component and aspect language designs. Can we develop a collection of component and aspect languages that can be plugged together in different ways for different applications? Can we use meta-level frameworks [21,22, 23, 24] to build such a collection?

What theoretical support can be developed for AOP? What kinds of theories can best describe the interaction between aspects and components and how they must be woven? Can such theories support development of a practical weaving toolkit?

What about the analysis and design process? What are good design principle for aspectual decomposition? What are good “module” structures for aspect programs? How can we train people to identify aspects? Clearly separate them? Write aspect programs? Debug AOP systems? Document AOP systems?

Another important area of exploration is the integration of AOP with existing approaches, methods, tools and development processes. As the examples in this paper show, AOP can be used as an improvement to existing techniques. To fulfill this promise it must be developed in a way that integrates well with those techniques.

8. RELATED WORK

In this section we give a brief survey of work related to ours. We start with work that is more closely related and proceed out to work that is less closely related.

8.1. WORK EXPLICITLY CONNECTED TO AOP

Several other groups have begun to explicitly consider their work in AOP terms. These include:

- Calton Pu et. al. at the Oregon Graduate Institute, in their work on Synthetix, are developing high performance, high portability and high adaptiveness OS kernels [3, 25]. In their work, the components are familiar functional elements of OS kernels. The aspects are primarily optimizations based on invariants that relate to how a service is being used. Their weaver technology uses partial evaluation to effectively specialize the kernel code for particular use cases. Their code is structured to expose as join points those places where an invariant becomes or ceases to be true.
- Karl Lieberherr et. al., at Northeastern University are developing techniques that make object-oriented programs more reusable and less brittle in the face of common program evolution tasks [2, 26, 27]. In their work, the component languages are existing OOPs like C++ and Java. Succinct traversal specifications [2] and context objects [27] provide aspect languages that can be used to address a variety of cross-cutting issues. Weaving of aspect programs that use succinct traversal specification is compile-time oriented, the join point representation is, roughly speaking, the class graph. Weaving of aspect pro-

grams that use context objects is more runtime oriented, the join points are the dynamic method and function calls.

- Mehmet Aksit et. al., at the University of Twente, have developed the composition filters object model, which provides control over messages received and sent by an object [1]. In their work, the component language is a traditional OOP, the composition filters mechanism provides an aspect language that can be used to control a number of aspects including synchronization and communication. Most of the weaving happens at runtime; the join points are the dynamic message sends and receives arriving at an object.

8.2. REFLECTION AND METAOBJECT PROTOCOLS

Aspect-oriented programming has a deep connection with work in computational reflection and metaobject protocols [28, 29, 30, 23, 31, 22]. A reflective system provides a base language and (one or more) meta-languages that provide control over the base language's semantics and implementation. The meta languages provide views of the computation that no one base language component could ever see, such as the entire execution stack, or all calls to objects of a given class. Thus, they cross-cut the base level computation. In AOP terms they are lower-level aspect languages whose join points are the "hooks" that the reflective system provides.

We have exploited this connection to great advantage in our work on AOP. When prototyping AOP systems we often start by developing simple metaobject protocols for the component language, and then prototype imperative aspect programs using them. Later, once we have a good sense of what the aspect programs need to do, we develop more explicit aspect language support for them.

The connection is particularly evident in section 6.1, where the aspect languages we provided could have been layered on top of a reflective architecture. Similarly, the loop fusion aspect described in Section 5.2 can be implemented, with some degree of efficiency, using the method combination facility in the CLOS metaobject protocol [30, 32]. This connection is also evident in the work mentioned in Section 8.1; both the Demeter work and the composition filters work have been described as being reflective facilities [33].

8.3. PROGRAM TRANSFORMATION

The goal of work in program transformation is similar to that of AOP. They want to be able to write correct programs in a higher-level language, and then mechanically transform those program into ones with identical behavior, but more efficient performance. In this style of programming, some of the properties the programmer wants to implement are written in an initial program. Other properties are added by passing that initial program through various transformation programs. This separation is similar in spirit to the component/aspect program separation.

But we believe that the precise notion of component and aspect are new to our AOP work, and provide us with additional value in designing the systems we have. It also appears that some transformations are aspectual in nature, but others are not. Transformation programs tend to operate in terms of the syntax of the program being transformed. If other join points are desired, it is the responsibility of the transformation program to somehow manifest them. Thus, while it is possible to layer some kinds

of aspect programs on top of a program transformation substrate, that is a separate piece of implementation work.

We would like to do a systematic analysis of the transformations developed by this community, to see which of them can be used for providing different kinds of aspect languages.

8.4. SUBJECTIVE PROGRAMMING

A natural question to ask is whether subjective programming is AOP or vice versa. We believe that AOP and subjective programming are different in important ways. Analogously to the way object oriented programming supports automatic selection among methods for the same message from different classes, subjective programming supports automatic combination of methods for a given message from different subjects. In both cases, the methods involved are components in the AOP sense, since they can be well localized in a generalized procedure. It is even possible to program in either an object-oriented style or a subjective style on top of an ordinary procedural language, without significant tangling. The same is not true of AOP. Thus, while the aspects of AOP tend to be about properties that affect the performance or semantics of components, the subjects of subjective programming tend to be additional features added onto other subjects. We believe that subjective programming is complementary to, and compatible with, AOP.

8.5. OTHER ENGINEERING DISCIPLINES

Many other engineering disciplines are based on well-established aspectual decompositions. For example, mechanical engineers use static, dynamic and thermal models of the same system as part of designing it. The differing models cross-cut each other in that the different properties of a system compose differently. Similarly, some software development tools explicitly support aspectual decomposition: tools for OMT [34, 35] methods let programmers draw different pictures of how objects should work.

9. CONCLUSIONS

We have traced the complexity in some existing code to a fundamental difference in the kinds of properties of a system that are being implemented. Components are units of a system's functional decomposition that can be cleanly captured in a generalized procedure. Aspects are properties that cannot be cleanly captured in a generalized procedure. We say that aspects cross-cut components.

Based on our analysis we have been able to develop programming language technology that supports clean abstraction and composition of both components and aspects. The key difference between AOP and other approaches is that AOP provides component and aspect languages with different abstraction and composition mechanisms. A special language processor called an aspect weaver is used to coordinate the co-composition of the two kinds of program.

We have had good success working with AOP in several testbed applications. The conceptual framework has helped us to design the systems, and the AOP-based implementations have proven to be easier to develop and maintain, while being comparably efficient to much more complex code written using traditional techniques.

ACKNOWLEDGEMENTS

Thanks to Karl Lieberherr, Carine Lucas, Gail Murphy and Bedir Tekinerdogan who generously provided extensive comments on earlier drafts of the paper, and to John Gilbert, Arthur Lee, Calton Pu, Alex Silverman and Marvin Theimer with whom we have had many discussions about AOP.

Thanks also to all the attendees of the AOP Friends Meetings, with whom we spent an enjoyable two days discussing AOP and related ideas: Mehmet Aksit, Lodewick Bergmans, Pierre Cointe, William Harrison, Jacques Malenfant, Satoshi Matsuoka, Kim Mens, Harold Ossher, Calton Pu, Ian Simmonds, Perri Tarr, Bedir Tekinerdogan, Mark Skipper, and Patrick Steyaert

BIBLIOGRAPHY

1. Aksit, M., *et al.* *Abstracting object interactions using composition filters.* in *ECOOP'93 Workshop on Object-Based Distributed Programming*. 1993: Springer-Verlag.
2. Lieberherr, K.J., I. Silva-Lepe, and C. Xaio, *Adaptive Object-Oriented Programming Using Graph-Based Customization*. *Communications of the ACM*, 1994. **37**(5): p. 94-101.
3. Pu, C., *et al.* *Optimistic Incremental Specialization: Streamlining a Commercial Operating System.* in *15th ACM Symposium on Operating Systems Principles (SOSP'95)*. 1995.
4. Parnas, D.L., *On the Criteria to be Used in decomposing Systems into Modules*. *Communications of the ACM*, 1972. **15**(2).
5. Parnas, D.L. *On a 'Buzzword': Hierarchical Structure.* in *IFIP Congress 74*. 1974: North-Holland Publishing Company.
6. Parnas, D.L. *Designing Software for Extension and Contraction.* in *3rd International Conference on Software Engineering*. 1978.
7. Mendhekar, A., G. Kiczales, and J. Lamping, *RG: A Case-Study for Aspect-Oriented Programming*, in *Submitted to OOPSLA 97 for possible publication*. 1997.
8. Mahoney, J.V., *Functional Visual Routines*, . 1995, Xerox Palo Alto Research Center: Palo Alto.
9. Kiczales, G., *et al.* *Open Implementation Design Guidelines.* in *International Conference on Software Engineering*. 1997. (Forthcoming).
10. Maeda, C., *et al.* *Open Implementation Analysis and Design.* in *Symposium on Software Reuse (To Appear, May 1997)*. 1997.
11. Kiczales, G., *Why are Black Boxes so Hard to Reuse?* 1994: Invited Talk, OOPSLA'94.
12. Kiczales, G., *Foil for the Workshop on Open Implementation*. 1994: Xerox PARC.

13. Henry, S. and D. Kafura, *Software Structure Metrics Based on Information Flow*. IEEE Transactions on Software Engineering, 1981. **SE-7**: p. 509--518.
14. McClure, C. *A Model for Program Complexity Analysis*. in *3rd International Conference on Software Engineering*. 1978. Los Alamitos, CA: IEEE Computer Society Press.
15. Yau, S. and J. Collofello, *Some Stability Measures for Software Maintenance*. tse, 1980. **SE-6**: p. 545--552.
16. Lopes, C.V. and G. Kiczales, *D: A Language Framework for Distributed Programming*, in *Submitted for Possible Publication, OOPSLA 97*. 1997.
17. Watanabe, T. and A. Yonezawa. *Reflection in an object-oriented concurrent language*. in *ACM Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 88)*. 1988. San Diego, CA.
18. Okamura, H., Y. Ishikawa, and M. Tokoro. *Al-1/d: A distributed programming system with multi-model reflection framework*. in *International Workshop on New Models for Software Architecture '92; Reflection and Meta-Level Architecture*. 1992.
19. Yokote, Y. *The Apertos Reflective Operating System: The Concept and its Implementation*. in *Conference on Object-Oriented Programming: Systems, Languages, and Applications*. 1992.
20. Ichisugi, Y., S. Matsuoka, and A. Yonezawa. *Rbc1: A reflective object-oriented concurrent language without a run-time kernel*. in *International Workshop on New Models for Software Architecture '92; Reflection and Meta-Level Architecture*. 1992.
21. Bobrow, D.G., et al., *Common Lisp Object System Specification*. SIGPLAN Notices, 1988. **23**.
22. Yonezawa, A. and T. Watanabe, *An Introduction to Object-Based Reflective Concurrent Computation*, in *Proceedings of the ACM SIGPLAN Workshop on Object-Based Concurrent Programming*, G. Agha, P. Wegner, and A. Yonezawa, Editors. 1989.
23. Matsuoka, S., T. Watanabe, and A. Yonezawa, *Hybrid Group Reflective Architecture for Object-Oriented Concurrent Reflective Programming*, in *European Conference on Object Oriented Programming*. 1991. p. 231-250.
24. Chiba, S. *A Metaobject Protocol for C++*. in *Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA 95)*. 1995. Austin: ACM.
25. Massalin, H. and C. Pu, *Threads and Input/Output in the Synthesis Kernel*, in *Proceedings of the 12th ACM Symposium on Operating Systems Principles*. 1989. p. 191-201.
26. Lopes, C.V. and K. Lieberherr. *Abstracting Process-to-Function Relations in Concurrent Object-Oriented Applications*. in *European Conference on Object-Oriented Programming*. 1994. Bologna, Italy: Springer-Verlag.

27. Seiter, L.M., J. Palsberg, and K.J. Lieberherr. *Evolution of Object Behavior Using Context Relations*. in *Fourth ACM SIGSOFT Symposium on the Foundations of Software Engineering*. 1996. San Francisco: ACM.
28. Smith, B.C., *Reflection and Semantics in a Procedural Language*
LCS Technical Report, . 1982, M.I.T.: Cambridge, MA.
29. Wand, M. and D.P. Friedman, *The Mystery of the Tower Revealed: A Non-Reflective Description of the Reflective Tower*, in *Proceedings of the ACM Conference on LISP and Functional Programming*. 1986, ACM. p. 298-307.
30. Kiczales, G., J.d. Riveres, and D.G. Bobrow, *The Art of the Metaobject Protocol*. 1991: MIT Press.
31. Okamura, H., Y. Ishikawa, and M. Tokoro. *Metalevel Decomposition in AL-1/D*. in *International Symposium on Object Technologies for Advanced Software*. 1993: Springer-Verlag.
32. Steele, G.L., *Common LISP: The Language, 2nd Edition*. 1990: Digital Press.
33. Lopes, C.V. and K. Lieberherr. *AP/S++: Case-Study of a MOP for Purposes of Software Evolution*. in *Reflection'96*. 1996. S. Francisco, CA.
34. Rumbaugh, J., *et al.*, *Object-Oriented Modeling and Design*. 1991: Prentice Hall.
35. Rational Software Corporation, <http://www.rational.com>, .