abc : An extensible AspectJ compiler*

Pavel Avgustinov[1], Aske Simon Christensen[2], Laurie Hendren[3], Sascha Kuzins[1], Jennifer Lhoták[3], Ondřej Lhoták[3], Oege de Moor[1], Damien Sereni[1], Ganesh Sittampalam[1], Julian Tibble[1]

[1] Programming Tools Group, Oxford University, United Kingdom
[2] BRICS, University of Aarhus, Denmark
[3] Sable Research Group, McGill University, Montreal, Canada

ABSTRACT

Research in the design of aspect-oriented programming languages requires a workbench that facilitates easy experimentation with new language features and implementation techniques. In particular, new features for AspectJ have been proposed that require extensions in many dimensions: syntax, type checking and code generation, as well as data flow and control flow analyses.

The AspectBench Compiler (*abc*) is an implementation of such a workbench. The base version of *abc* implements the full AspectJ language. Its frontend is built, using the Polyglot framework, as a modular extension of the Java language. The use of Polyglot gives flexibility of syntax and type checking. The backend is built using the Soot framework, to give modular code generation and analyses.

In this paper, we outline the design of *abc*, focusing mostly on how the design supports extensibility. We then provide a general overview of how to use *abc* to implement an extension. Finally, we illustrate the extension mechanisms of *abc* through a number of small, but non-trivial, examples. *abc* is freely available under the GNU LGPL.

1. INTRODUCTION AND MOTIVATION

The design and implementation of aspect-oriented programming languages is a buoyant field, with many new language features being developed. In the first instance, such features can be prototyped in a system like the Aspect Sand Box [11] via a definitional interpreter. Such interpreters are useful in defining the semantics and in explaining the compilation strategy of new language features [22]. The acid test for new language features is, however, their integration into a full, industrial-strength language like AspectJ. That requires a highly flexible implementation of AspectJ, to serve as a workbench for aspect-oriented programming language research.

The purpose of this paper is to present *abc*, the AspectBench Compiler for AspectJ, which supports the whole of the AspectJ language implemented by *ajc* 1.2, and which has been specifically designed to be an extensible framework for implementing AspectJ extensions. *abc* is freely available under the GNU LGPL and can be downloaded from [1].

Challenges. An AspectJ compiler is already a complex piece of software, which, in addition to the normal frontend and backend components of a compiler, must also support a matcher (for name patterns) and a weaver (both for intertype declarations and for advice). Furthermore, the kinds of extensions that have been suggested for AspectJ vary from fairly simple pointcut language extensions to more complex concepts which require modifications in the type system, matcher and weaver. To make the challenges explicit, we briefly review some previous work by others that has motivated our design.

At one end of the spectrum, there are fairly small extensions that require changes primarily to the syntax. An example of this kind is the *name pattern scopes* proposed by Colyer and Clement [8], which provide an abstraction mechanism for name patterns. To support this type of extension, our workbench needs an easy way of extending the syntax, as well as introducing named patterns into the environment.

A more involved extension is the *parameteric introductions* of Hanenberg and Unland [15]. These are intertype declarations that depend on parameters evaluated at weave-time. Their integration into AspectJ requires substantial changes to the type system as well as the intertype weaver. This kind of extension thus motivates a highly flexible implementation of types.

Most proposals for new features in AspectJ are, however, concerned with the dynamic join point model. In [24], Sakurai *et al.* propose *association aspects*. These provide a generalisation of per-object instantiation, where aspect instances are tied to a group of objects to express behavioural relationships more directly. This requires not only changes to the frontend, but also substantial changes to code generation. Making such code generation painless is another design goal of our workbench.

The community as a whole is concerned with finding ways of singling out join points based on semantic properties rather than naming. For instance, Kiczales has proposed a new type of pointcut, called *predicted cflow* [19]. **pcflow**(p) matches at a join point if there may exist a path to another join point where p matches. It is correct to let **pcflow**(p) match everywhere, but that would lead to inefficient programs. An efficient implementation of **pcflow**(p) needs substantial, interprocedural program analysis. Our workbench needs to provide a framework for building such analyses.

In fact, examples where efficient implementation necessitates an analysis framework abound. Particular instances include the *data flow pointcuts* of Masuhara and Kawauchi [21], and the *trace-based aspects* of Douence *et al.* [10], as well as the *communication history*

*This work was supported, in part, by NSERC in Canada and EPSRC in the United Kingdom. Our thanks to Chris Allan for his comments on a draft of this paper.

AOSD'05, March 14-18, 2005, Chicago, USA

aspects of Walker and Viggers [28].

All of the above are additions to the AspectJ language, but, of course, restrictions can be equally important in language design. One promising example is the proposal of Aldrich to restrict the visibility of join points to those that are explicit in the interface of a class [2]. We aim to support the implementation of such restrictions, and this requires a flexible implementation of the type system and the pointcut matcher.

Finally, we note that the implementation of advanced static checking tools for aspect-oriented programs, such as those investigated by Krishnamurthi *et al.* [20], require all types of extension discussed above, ranging from simple variations in syntax to making advanced analyses such as escape analysis take into account the effects of advice.

In summary, we can see that an extensible AspectJ compiler must be able to handle a wide variety of extensions, possibly touching on many components of the compiler, including the frontend scanner and parser, the type checker, the matcher and weaver, and potentially requiring relatively sophisticated program analysis to ensure correctness and efficiency.

Design Goals. One approach to implementing a language extension is to modify an existing compiler. However, this is not always the best approach, since existing compilers may not have been designed with extensiblity as one of the main goals. Furthermore, they may be constrained to work with infrastructures which themselves are not easily extensible. In the case of AspectJ, the only pre-existing implementation is *ajc*, which is designed to support fast and incremental compilation and also to interact closely with the Eclipse toolset.

Our approach was to design and implement *abc*, the AspectBench Compiler, with extensibility as its primary design goal (we also aimed for an optimising implementation of AspectJ, but that is outside the scope of this paper). To support extensibility, we distilled the following requirements from the above discussion of the challenges involved.

simplicity: It must be relatively simple to develop new extensions. Users of the framework should not need to understand complicated new concepts or a complex software design in order to implement their extensions.

modularity: We require two kinds of modularity. First, the compiler workbench itself should be very modular, so that the different facets of each extension can be easily identified with the correct module of the workbench.

Second, the extension should be modular (separate from the workbench code). Users of the workbench should not need to touch existing code; rather, they should be able to describe the extensions as specifications or code that is separate from the main code base.

proportionality: Small extensions should require a small amount of work and code. There should not be a large overhead required to specify an extension.

analysis capability: The compiler workbench infrastructure should provide both an intermediate representation and a program analysis framework. This is necessary for two reasons. First, some extensions may require relatively sophisticated analyses to correctly implement their semantic checks and weaving. Second, some extensions may lead to a lot of runtime overhead unless compiler optimisation techniques are used to minimise that overhead.

The abc approach. To meet these objectives, we decided to build on existing, proven tools, namely the Polyglot extensible compiler framework for the frontend [23], and the Soot analysis and transformation framework for the backend [26]. Indeed, Polyglot has been shown to meet the criteria of simplicity, modularity and proportionality on a wide variety of extensions to the syntax and type system of Java. By the same token, Soot has been shown to meet all the above criteria for code generation, analysis and optimisation.

Given the success of these building blocks, we felt it extremely important to design *abc* so that both are used *as is*, without any changes that are specific to *abc*. As explained in Section 2 below, this has dictated an architecture where the frontend separates the AspectJ program into a pure Java part and a part containing instructions for the weaver.

Contributions. The contributions of this paper are the following:

- We have identified the requirements for a workbench for research in aspect-oriented programming languages by analysing previous research in this area.
- We present *abc*, an instance of such a workbench with a clean, extensible architecture.
- We have validated our architecture against these requirements with a number of small but non-trivial examples.
- The extensibility of Polyglot can be seen as a form of aspect-orientation, and thus *abc* is itself a substantial exercise in aspect-oriented software development, with the primary goal of disentangling new, experimental features from the existing base compiler.

Paper Structure. The structure of this paper is as follows. In Section 2, we first give an overview of the main building blocks of *abc*, namely Polyglot and Soot, and show their role in the overall architecture of *abc*. Next, in Section 3 we sketch the main points of extensibility in *abc*. We then turn to describe some modest but representative examples of AspectJ extensions in Section 4, and their implementation in Section 5. The approach of *abc* is contrasted with existing work in Section 6. Finally, in Section 7 we draw some conclusions from our experience in building *abc*, and we explore possible directions for future research.

2. ARCHITECTURE

As stated in the introduction, *abc* is based on the Polyglot extensible compiler framework [23] and the Soot bytecode analysis and transformation framework [26]. Using Polyglot as an extensible frontend enables customisation of the grammar and semantic analysis; in the backend, Soot provides a convenient intermediate representation on which to implement the weaving of extensions, as well as tools for writing any program analyses that extensions may require. Polyglot can read class files to process library code; Soot can also read in class files, and indeed *abc* can weave into jar files in the same way as the *ajc* compiler for AspectJ.

Because *abc* works with an unmodified Soot and Polyglot, it is easy for us, as the developers of *abc* itself, to update to the latest versions of Soot and Polyglot as they are released. By the same token, authors of AspectJ extensions can upgrade to new versions of *abc* without difficulty. This independence was achieved mainly by separating the AspectJ-specific features in the code being processed

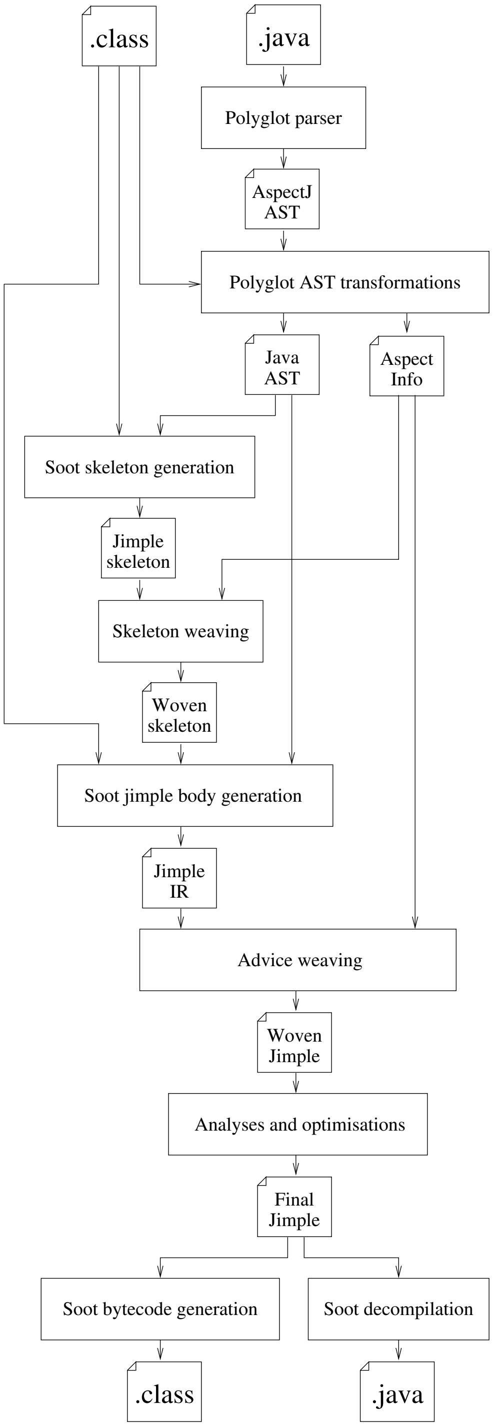

Figure 1: *abc* overall design

from standard Java code. In the frontend, *abc* generates a plain Java abstract syntax tree (AST) and a separate aspect information structure containing the aspect specific information. We call the aspect information structure the *AspectInfo*. The unmodified backend can read in the AST (because it is plain Java), and *abc* then uses the *AspectInfo* to perform all required weaving. A simplified diagram of the architecture of *abc* is shown in Figure 1.

In the following subsections, we describe Polyglot and Soot in the context of *abc*, with a focus on how they contribute to extensibility. Finally we discuss in some more detail how the two parts are connected.

2.1 Polyglot

Polyglot [23] is a frontend for Java intended for implementing extensions to the base language. In its original configuration, Polyglot first parses Java source code into an abstract syntax tree (AST), then performs all the static checks required by the Java language in a number of passes which rewrite the tree. The output of Polyglot is a Java AST annotated with type information, which is written back to a Java source file. Polyglot is intended to perform all compile-time checks; when a class has passed through all of the passes in Polyglot, the resulting Java file should be compilable without errors by any standard Java compiler. When Polyglot is used as a frontend for Soot, the *Java to Jimple* module inside Soot compiles the final AST into the Jimple intermediate representation instead of writing it out to a Java file. Therefore, in *abc*, the final Polyglot passes separate the AspectJ program into pure Java (which is passed to the Java to Jimple module in Soot) and instructions for the weaver.

Several features of Polyglot make it well-suited for writing extensions, and also help to make those extensions themselves extensible. Polyglot allows a new grammar to be specified as a collection of modifications to an existing grammar, where these modifications are given in a separate specification file, not in the original grammar file. The AspectJ grammar we developed for *abc* is specified as an extension of the Java grammar, and the grammars for extensions are in turn specified as modifications to the AspectJ grammar.

Polyglot makes heavy use of interfaces and factories, making it easy to extend or replace most of its parts, such as the type system or the symbol table, as well as the list of rewrite passes that are performed on the AST. Each pass in Polyglot non-destructively rewrites the input tree. As a result, it is easy to insert new passes in between existing ones, and each pass typically performs only a small amount of work compared to traditional compiler passes. In *abc*, we have added many AspectJ-specific passes, and it is easy for extensions to add further passes of their own.

Each AST node in Polyglot uses a mechanism of *extensions* and *delegates* to allow methods to be replaced or added in the middle of the existing class hierarchy, achieving an effect similar to what can be done in AspectJ using intertype declarations, but in plain Java. This mechanism is commonly used by extensions of *abc* to modify existing AST nodes.

2.2 Soot

Soot [26], which is used as the back-end of *abc*, is a framework for analysing and transforming Java bytecode. The most important advantage of using Soot as the backend, both for developing *abc* itself and for extending the language, is Jimple, Soot's intermediate representation. Soot provides modules to convert between Jimple, Java bytecode, and Java source code. It furthermore includes implementations of standard compiler optimisations, which *abc* applies after weaving. We have already observed significant speedups from these optimisations alone (preliminary results are accessible from [1]). In addition to already implemented analyses and transformations, Soot has tools for writing new ones, such as control flow graph builders, definition/use chains, a fixed-point flow analysis framework, and a method inliner. These features are useful for implementing extensions that need to be aware of the intra-procedural behaviour of the program, such as pointcuts describing specific points in the control flow graph.

The Jimple intermediate representation is a typed, stack-less, three-address code. Rather than representing computations with an implicit stack, each Jimple instruction explicitly manipulates specified local variables. This representation simplifies weaving of advice, both for standard AspectJ features and for extensions. If it were weaving into bytecode directly, the weaver would need to consider the effect of the woven code on the implicit execution stack, and generate additional code to fix up the stack contents. None of this is necessary when weaving into Jimple. Moreover, when values from the shadow point are needed as parameters to the advice, they are readily available in local variables; the weaver does not have to sift through the computation stack to find them.

As input, Soot can handle both class files and Java source files. To convert bytecode to Jimple, Soot introduces a local variable to explicitly represent each stack location, splits the variables to separate independent uses of the same location, and infers a type [13] for each variable. To convert source code to Jimple, Soot first uses Polyglot to construct an AST with type information, and then generates Jimple code from the AST. This process does not need to be modified in *abc*, because *abc* passes Soot a plain Java AST, keeping all the aspect-specific information in the separate aspect information structure. Normally, after all processing, Soot converts the Jimple code into bytecode and writes it to class files, but it also includes a decompiler which is very useful for viewing the effects of aspects and AspectJ extensions on the generated code.

2.3 Connecting Polyglot and Soot

We conclude the discussion of *abc*'s architecture by examining in closer detail how Polyglot and Soot interact. A key component of this interaction is the separation of the AspectJ AST into a pure Java AST and the auxilliary *AspectInfo* structure. This transformation enables *abc* to use the existing facility in Soot for translating a Polyglot AST into the Jimple IR.

The Java AST is basically the AspectJ program with all AspectJ-specific language constructs removed. The *AspectInfo* structure contains complete information about these constructs. In cases where these contain actual Java code (advice bodies, **if** pointcut conditions, intertype method/constructor bodies, intertype field initialisers), the code is placed in placeholder methods in the Java AST.

The Java AST only contains Java constructs, but it is incomplete in the sense that it may refer to class members which do not exist or are not accessible in the unwoven Java program. More specifically, the Java AST will in general not be compilable until all *declare parents* and intertype declarations have been woven into the program. The first of these can alter the inheritance hierarchy, and the second can introduce new members that the pure Java parts may refer to. Since both of these features may be applied to class files (for which we do not have an AST representation), it is not possible to perform this part of the weaving process on the Polyglot representation before passing the AST to Soot.

Fortunately, Soot allows us to conduct the conversion from Java to Jimple in two stages, and the application of *declare parents* and intertype weaving can happen in between. In the first stage, Soot builds a class hierarchy with mere stubs for the methods: it is a skeleton of a full program in Jimple, without method bodies. In the second stage, Soot fills in method bodies, either by converting bytecode from class files, or by compiling AST nodes.

This setup permits both static weaving and advice weaving to work on the Jimple IR, largely independent of whether the Jimple code was generated from source code or bytecode. And since the skeleton that is filled out in the second stage has the updated hierarchy and contains all intertype declarations, all member references in the code are resolved correctly in the translation into Jimple.

3. DEFINING AN EXTENSION

We now outline the basic steps needed to create an extension, in a general manner. This description is intended to give the reader an impression of the extension mechanisms available in *abc*, without delving into excessive detail. After this generic description, we shall introduce some concrete examples in Section 4, and show how the basic steps are instantiated in Section 5.

This section serves two purposes. First, to outline how we build on the existing extension mechanisms of Polyglot and Soot to achieve extensibility in *abc* (Sections 3.2, 3.3, 3.4, and 3.8). Second, we wish to present some design decisions that are unique to *abc*, which address specific issues regarding the extension of AspectJ (Sections 3.1, 3.5, 3.6, 3.7 and 3.9).

3.1 Syntax

The first step in implementing a new extension is usually defining what additional syntax it will introduce to the language. Making *abc* recognise the extended language involves changing the lexer and the parser that it uses. Polyglot already handles extending grammars in a very clean and modular fashion, however the standard Polyglot lexer is not extensible — extensions are expected to create their own lexer by copying it and making appropriate modifications. Thus, in this subsection we describe our approach to making an extensible lexer in some detail, and then briefly summarise the Polyglot mechanism for extending grammars.

Lexer. We have designed the lexer for *abc* to support a limited form of extensibility that has been sufficient for the extensions we have written so far. Specifically, the set of keywords recognised by the lexer can be modified by an extension, and the actions taken by the lexer when encountering one of these keywords are customisable. More complex extensions can still be achieved by reverting to Polyglot's approach of copying and modifying the lexer definition. This is in agreement with the principle of proportionality which was stated as a design goal — small extensions are easy, and complex ones are possible. It is a topic of future work to improve the extensibility, perhaps by specifying the lexer as a parser in itself.

The lexical analysis of AspectJ is complicated by the fact that there are really several different languages being parsed: ordinary Java code, aspect definitions, and pointcut definitions. Consequently, the *abc* lexer is stateful — it recognises different tokens in different contexts. The following example illustrates one kind of problem that is dealt with by the introduction of lexer states:

*if**.*1.*Foo*+.*new*(..)

The expected interpretation of such a string as Java code and as part of a pointcut will be very different; for example, in Java, we would expect "1." to become a floating point literal, whereas in the pointcut language the decimal point would be viewed as a dot separating elements of a name pattern. Similarly, "*" in Java should be scanned as an operator, while in pointcuts, it is part of a name pattern. Note also the use of what would be keywords in Java mode (**if** and **new**) as part of a pattern.

An important part of designing a stateful lexer is specifying when the lexer should switch to a different state without adding too much complexity. The general pattern we use is to maintain a stack of states, and recognise the end of a state when we reach an appropriate closing bracket character for that state. For example, normal Java code is terminated by the '}' character. Of course, braces can be nested, so we need to recognise opening braces and also count the nesting level. For more details regarding the lexer states in *abc*, see section 5.1.

File **X**	File **Y**
S ::= a | b | c	*include* **X** *extend* S ::= d | e
File Z	**Result**
include **Y** *drop* S ::= b | d	S ::= a | c | e

Figure 2: Grammar extension mechanism

Parser. The *abc* parser is generated by PPG [5], the LALR parser generator for extensible grammars which is included in Polyglot [23]. PPG allows changes to an existing grammar to be entered in a separate file, overriding, inheriting and extending productions from the base grammar. This results in modular extensions, which can easily be maintained should the base grammar change.

The example in Figure 2 (using simplified non-PPG syntax) demonstrates the basic principles. An existing grammar can be imported with the "*include*" keyword. New production rules can then be specified, and one can change existing rules using the keywords "*extend*" and "*drop*" to add and remove parts of the rule. More advanced changes, such as modifying the precedence of operators, are also possible. For further details on the specification of grammar, see [5].

3.2 Type system

Polyglot provides convenient facilities for extending the type system. As a minimum, this involves introducing a new kind of type object and lookup functions for these new entitities in the environment. The new type of environment is then invoked by overriding the environment factory method in a subclass of *AJTypeSystem*, which describes the type system of AspectJ itself.

To illustrate, consider the introduction of named class pattern expressions [8]. We would need to introduce a new type object to represent such names, say *NamedCPEInstance* (in Polyglot, it is convention that identifiers for type classes end with ... *Instance*). The environment then maps (possibly qualified) names to objects of type *NamedCPEInstance*.

The semantic checks for named patterns must enforce the requirement that there be no cycles in definitions, since recursively defined named patterns do not make sense. A similar check has already been implemented for named pointcuts, and it involves building a dependency graph. Such data structures necessary for semantic checks are typically stored in the type objects (here *NamedCPEInstance*): because Polyglot operates by rewriting the original tree, it is not possible to store references to AST nodes.

Examples such as the parametric introductions of Hanenberg and Unland [15] would require more invasive changes in the type system, for example by subclassing *InterTypeMethodInstance* (the signature of a method introduced via an intertype declaration) to take account of the parameters that are to be evaluated at compile-time.

3.3 Semantic checks

New semantic checks are usually implemented by overriding the appropriate method on the relevant AST nodes. The most obvious place for simple checks is in the *TypeChecker* pass; every AST node implements a *typeCheck*(*TypeChecker*) method. The type checker is run after all variable references are resolved; all checks that do not require further data structures are typically put in the *typeCheck* method.

Later passes use data flow information to check initialisation of local variables and the existence of **return** statements. Again, each AST node implements methods to build the control flow graph for these purposes. In the base AspectJ implementation, these are, for example, overridden to take into account the initialisation of the result parameter in **after returning** advice, and extensions can make variations of their own.

AspectJ is somewhat unusual in that some semantic checks have to be deferred to the weaver. For example, it is necessary to type check the results of **around** advice at each point where it is woven in. Because *abc* maintains precise position information throughout the compilation process, such errors can still be pinpointed to the appropriate locations in the source.

3.4 Rewriting

The normal use of Polyglot is as a source-to-source compiler for extensions to Java, where the final rewriting passes transform new features into an equivalent pure Java AST. *abc* is different in that most of the transformation happens at a later stage, when weaving into Jimple. It is, however, often useful to employ Polyglot's original paradigm when implementing extensions to AspectJ that have an obvious counterpart in AspectJ itself.

For example, consider again the feature of named class pattern expressions. A simple implementation would be to just inline these after appropriate semantic checks have been done, so that nothing else needs to change in the compiler. Such inlining would be implemented as two separate passes, one to collect the named pattern definitions and the other to inline them — the two would then communicate via an explicit data structure that is common to both passes. As said, it is not recommended to store pieces of AST explicitly unless they are immediately transformed away.

abc does extensive rewriting of the tree prior to conversion to Jimple. This consists of introducting new placeholder methods (for instance for advice bodies), and storing instructions for the weaver in the *AspectInfo*. Extensions can participate in this process by implementing methods that are called by the relevant passes.

3.5 Join points

Introducing new pointcuts will often involve extending the set of possible join points. For example, implementation of a pointcut that matches when a cast instruction occurs would require the addition of a join point at such instructions.

Many new join points will follow the pattern of most existing AspectJ join points and apply at a single Jimple statement. These can be added by defining a new factory class that can recognise the relevant statements, and registering it with the global list of join point types.

For more complicated join points, it will be necessary to override the code that iterates through an entire method body looking for join point shadows. The overriding code can do any required analysis of the method body to find instances of the new join points (for example, one might want to inspect all control flow edges to find the back edges of loops [17]), and then call the original code to find all the "normal" join point shadows.

3.6 Pointcuts

As pointed out in the introduction, there are many proposals for new forms of pointcuts in AspectJ. To meet our objective of proportionality (small extensions require little work), we have designed an intermediate representation of pointcuts that is more regular than the existing pointcut language of AspectJ. This makes it easier to compile new pointcut primitives to existing ones.

Specifically, the backend pointcut language partitions pointcuts into the four categories listed below. Some of the standard AspectJ pointcuts fit directly into one of these categories and are simply duplicated in the backend, while others are must be transformed from AspectJ into the representation used in *abc*.

- Lexical pointcuts are restrictions on the lexical position of where a pointcut can match. For example **within** and **withincode** fall into this category.
- Shadow pointcuts pick out a specific join point shadow within a method body. The **set** pointcut is an example.
- Dynamic pointcuts match based on the type or value of some runtime value. Pointcuts such as **if**, **cflow** and **this** are of this kind.
- Compound pointcuts represent logical connectives such as &&.

The motivation for this categorisation is that it allows the implementation of each backend pointcut to be simpler and more understandable, which in particular makes it easier for extension authors to define new pointcuts.

An example of an AspectJ pointcut that does not fit into this model directly is the **execution**(⟨*MethodPattern*⟩) pointcut, which specifies both that we are inside a method or constructor matching *MethodPattern* and that we are at the execution join point. The backend pointcut language therefore views this as the conjunction of a lexical pointcut and a shadow pointcut.

To add a new pointcut, one or more classes should be added to the backend, and the frontend AST nodes should construct the appropriate backend objects during the generation of the *AspectInfo* structure.

The backend classes are responsible for deciding whether or not the pointcut matches at a specific location. If this cannot be statically determined, then the pointcut should produce a *dynamic residue* which can generate the required runtime code.

3.7 Advice

It appears that there are few proposals for truly novel types of advice: most new proposals can be easily rewritten to the existing idioms of before, after and around. For example, the proposal for "tracecuts" in [10] reduces to a normal aspect, where a state variable tracks the current matching state, and each pattern/advice pair translates into after advice. Such new types of advice are thus implemented via rewriting, in the standard paradigm of Polyglot.

Nonetheless, adding a new kind of advice that follows the AspectJ model of advice is straightforward: simply implement a new class and define how code should be generated to call that piece of advice and where in the join point shadow this code should go. For example, the bookkeeping required for **cflow** is implemented as a special kind of advice that weaves instructions both at the beginning and end of a shadow.

3.8 Optimisations

The straightforward implementation of a new extension may result in inefficient runtime code. Even in the basic AspectJ language, there are a number of features that incur significant runtime penalties by default, but in many cases can be optimised. *abc* aims to make it as easy as possible to implement new optimisations, whether for the base language or for extensions. In particular, it is straightforward to transform the AST in the frontend and the Jimple intermediate code in the backend.

Taking an example from the base AspectJ language, construction of the **thisJoinPoint** is expensive because it must be done each time a join point is encountered at runtime. *abc* (like *ajc*) employs two strategies for mitigating this overhead. Firstly, some advice bodies only ever make use of the *StaticPart* member of **thisJoinPoint**, which only needs to be constructed once. A Polyglot pass in the frontend is used to identify advice bodies where this is the case and transform the uses to **thisJoinPointStaticPart** instead.

Secondly, the runtime code generated delays construction until as late as possible in case it turns out not to be needed at all; this is complicated by the fact that *if* pointcuts as well as advice bodies may make use of it, so construction cannot simply be delayed until the advice body runs. *abc* generates code that instantiates the **thisJoinPoint** variable where needed *if it has not already been instantiated*, using *null* as a placeholder until that point. The Jimple code is then transformed to remove unnecessary checks and initialisations, using a variation of Soot's intraprocedural nullness analysis which has special knowledge that the **thisJoinPoint** factory method cannot return *null*.

3.9 Runtime library

The runtime library for AspectJ serves two purposes. Firstly, it contains bookkeeping classes necessary for the implementation of language constructs such as **cflow**. Extensions such as data flow pointcuts [21] would require a similar runtime class in order to store dynamic data about the source of the value in a particular variable.

Secondly, the runtime provides the objects accessible through the *thisJoinPoint* family of special variables; these make information about the current join point available to the programmer via reflection. Any new pointcut introduced is likely to have unique signature information which would be accessible to the user via an extension of the *Signature* interface. For example, the standard AspectJ runtime contains, amongst others, *AdviceSignature*, *FieldSignature*, and *MethodSignature*.

4. EAJ — AN ASPECTJ EXTENSION

This section describes a few particular extensions to the AspectJ language that we have implemented. These extensions have been chosen to illustrate the most salient of the mechanisms that were described in the previous section. The full source code for these examples is included with the standard distribution of *abc* [1]. For ease of reference, the extended language is named *eaj*; one compiles *eaj* programs with the command '*abc -ext abc.eaj*'. This is the usual way of invoking extensions with *abc*.

4.1 Private pointcut variables

In AspectJ, the only way to introduce new variables into a pointcut is to make them explicit parameters to a named pointcut definition or advice. It is sometimes convenient, however, to simply declare new variables whose scope is only part of a pointcut expression, without polluting the interface of the pointcut. For example, it might be desired to check that the value of an argument being passed has certain properties, without actually using that value in the advice body. The new keyword **private** introduces a locally scoped pointcut variable. For instance, the following pointcut could be used to check that the argument is either a negative *int* or a negative *double*:

```
pointcut negativefirstarg() :
        private (int x) (args(x)  &&  if(x < 0))
     || private (double x) (args(x)  &&  if(x < 0));
```

4.2 Global pointcuts

It is very common for many pieces of advice to share a common conjunct in their pointcut. The idea of a *global* pointcut is to write these common conjuncts only once. An example use is to restrict the applicability of every piece of advice within a certain set of aspects. For example, we might write:

global : * : !**within**(*Hidden*);

This would ensure that no advice within any aspect could apply within the *Hidden* class.

As another example, it is often useful to prevent advice from an aspect applying within that aspect itself. The following declaration (for aspect *Aspect*) can achieve this more concisely than putting the restriction on each piece of advice:

global : *Aspect* : !**within**(*Aspect*);

In general, a global pointcut declaration can be put anywhere a named pointcut declaration can be (*i.e.* directly within a class or aspect body). The location of such a declaration has no effect on its applicability, except that name patterns within such a declaration will only match classes and aspects visible from the scope of that declaration.

The general form of a global pointcut declaration is as follows:

global : ⟨*ClassPattern*⟩ : ⟨*Pointcut*⟩ ;

It has the effect of replacing the pointcut of each advice declaration in each aspect whose name matches *ClassPattern* with the conjunction of the original pointcut and the global *Pointcut*.

4.3 Cast pointcuts

The purpose of the *cast* pointcut is to match whenever a value is cast to another type. A corresponding new type of join point shadow is added which occurs at every cast instruction, whether for reference or primitive types, in the bytecode of a program.

To illustrate, the following piece of advice can be used to detect runtime loss of precision caused by casts from an *int* to a *short*:

```
before(int i):
        cast(short) && args(i)
    && if(i < Short.MIN_VALUE
        || i > Short.MAX_VALUE)
{
    System.err.println("Warning: loss of " +
                        " precision casting " +
                        i + " to a short.");
}
```

In general the syntax of a **cast** pointcut is **cast**(⟨*TypePattern*⟩); this will match at any join point where the static result type of the cast is matched by *TypePattern*. In keeping with the pattern of other primitive pointcuts, the value being cast from can be matched by the **args** pointcut, and the result of the cast can be matched by the optional parameter to **after returning** advice (and is returned by the **proceed** call in **around** advice).

4.4 Throw pointcuts

The **throw** pointcut is introduced in the developer documentation for *ajc* [18], and we have implemented it in *eaj* to compare the ease-of-extension of both compilers. It matches a new join point shadow which occurs at each throw instruction.

The following example demonstrates how extended debugging information can be produced in the event of a runtime exception, using a piece of advice:

```
before(Debuggable d):
    this(d) && throw() && args(RuntimeException)
{
    d.dumpState();
}
```

5. IMPLEMENTING EAJ USING ABC

We have given a broad outline of how extensions are constructed and discussed some specific extensions that we have implemented. We now show in detail how this was done, both to provide a guide for others and to enable a realistic assessment of the work involved.

The starting point to extending *abc* is the *AbcExtension* class. An extension can be specified at runtime by passing its core package name to *abc* with the *-ext* flag; the *AbcExtension* class from this package is then loaded by reflection; all the extensibility hooks in *abc* are passed through this class. There is a default implementation of this class in the *abc.main* package, which extensions must subclass.

Another key class is *ExtensionInfo*. This is part of the extensibility mechanism of Polyglot; all frontend extensions (except for the lexer) are registered by subclassing this class. New instances of this class are returned by the subclassed *AbcExtension*.

5.1 Extending the lexer

As described in Section 3.1, *abc*'s lexer is stateful. There are four main lexer states for dealing with the different sub-languages of AspectJ: JAVA, ASPECTJ, POINTCUT and POINTCUTIFEXPR. The first three are used in the obvious contexts. The POINTCUTIFEXPR state must be separate from the normal JAVA state because the **if** pointcut allows a Java expression to be nested inside a POINTCUT, but whereas the JAVA state is terminated by a '}', we need to return to the POINTCUT state when reaching a matching closing ')' character.

Keywords for each state are stored in state-specific *HashMap*s which map each keyword to an object implementing the *LexerAction* interface. This interface declares a method

```
public int getToken(AbcLexer lexer)
```

which is called when the corresponding keyword is recognised. Its return value is turned into a parser token and passed to the parser for further analysis. A reference to the lexer instance is passed as a parameter to *getToken*(...), so that side effects that affect the lexer (like changing the lexer state) are possible. A default implementation of this interface is supplied, which offers sufficient functionality to associate keywords with parser tokens and (optionally) change the lexer state; custom implementations of *LexerAction* can provide more flexibility. Note that the default implementation provides functionality sufficient for all but 5 (out of more than 90) Java and AspectJ keywords.

Implementing the *eaj* extensions required adding several new keywords. In particular, "*cast*" was introduced as a keyword in the POINTCUT state, and "*global*" as a keyword in all four lexer states. Note that "*private*" and "*throw*" are already keywords in all states, and so do not need to be introduced specifically for the private pointcut variables and throw pointcut extensions. Here is the code that adds the keywords to the respective states:

```
public void initLexerKeywords(AbcLexer lexer)
{
    // keyword for the "cast" pointcut extension
    lexer.addPointcutKeyword("cast",
        new LexerAction_c(new Integer
            (abc.eaj.parse.sym.PC_CAST)));

    // keyword for the "global pointcut" extension
    lexer.addGlobalKeyword("global",
        new LexerAction_c(new Integer
            (abc.eaj.parse.sym.GLOBAL),
                new Integer(lexer.pointcut_state())));

    // Add the base keywords
    super.initLexerKeywords(lexer);
}
```

Note that both keywords use the default implementation of *LexerAction*, i.e. the *LexerAction_c* class. We see the one-argument and two-argument constructors for that class. The first argument is always the parser token that should be returned for the keyword, the second argument (if present) is the lexer state that should be selected after the keyword. As stated above, further logic can be implemented by subclassing *LexerAction_c*.

5.2 Extending the parser

The grammar fragment below shows how two new productions are added for private pointcut variables and the cast pointcut, which can appear anywhere a normal pointcut could:

```
extend basic_pointcut_expr ::=
    PRIVATE:x LPAREN formal_parameter_list_opt:a RPAREN
    LPAREN pointcut_expr:b RPAREN:y
    {:
        RESULT =
            parser.nf.PCLocalVars(parser.pos(x,y), a, b);
    :}
|   PC_CAST:x LPAREN type_pattern_expr:a RPAREN:y
    {:
        RESULT =
            parser.nf.PCCast(parser.pos(x,y), a);
    :}
;
```

The fragment closely resembles code one would use with the popular CUP parser generator, apart from the **extend** keyword, which signifies that these two productions are to be added to the rules that already exist for *basic_pointcut_expr*. PPG also allows you to drop productions, transfer productions from one non-terminal to another, and override the productions of a particular non-terminal; it is described in detail in [5].

5.3 Adding new AST nodes

As mentioned above, *abc*'s frontend is built on the Polyglot extensible compiler framework [23]. In fact, from Polyglot's point of view, *abc* is just another extension. This means that *abc* "inherits" all the extensibility mechanisms provided by Polyglot.

In particular, adding new AST nodes is common when writing compiler extensions, and thus it is important to provide an easy and robust mechanism for doing so.

All four extensions discussed above required new AST nodes. For the sake of brevity we will only present the node introduced by the global pointcut extension here — the other cases are handled very similarly.

In order to write a clean Polyglot extension, one has to adhere to the rigorous use of factories and interfaces to create nodes and invoke their members, respectively. The first step is therefore to define an interface for the new AST node, declaring any functionality it wants to present to the outside world:

```
public interface GlobalPointcutDecl extends PointcutDecl
{
    public void registerGlobalPointcut(GlobalPointcuts visitor,
                                       Context context,
                                       EAJNodeFactory nf);
}
```

We provide a method to insert the pointcut into a static data structure keeping track of the global pointcuts defined in the program (*cf.* Section 5.4). Note that the interface extends *abc*'s *PointcutDecl* interface, so it provides all the functions relevant to a pointcut declaration.

The next step is to write the class implementing that interface. Some boilerplate code is required (a constructor and methods to allow visitors to visit the node), and, of course, the method *registerGlobalPointcut*() is given a concrete implementation.

In order to make sure we can instantiate this new node type, we subclass *abc*'s default node factory (which, in turn, is derived from Polyglot's node factory) and create a method for obtaining an instance of *GlobalPointcutDecl*:

```
public GlobalPointcutDecl
        GlobalPointcutDecl (
                Position pos,
                ClassnamePatternExpr aspect_pattern,
                Pointcut pc, String name,
                TypeNode voidn )
{
    return new GlobalPointcutDecl_c(pos, aspect_pattern,
                                    pc, name, voidn);
}
```

Now the extended parser can produce *GlobalPointcutDecl* objects when it encounters the appropriate tokens (*cf.* listing in Section 5.2).

Note that all changes are local to new classes we created (in fact, these classes are in a completely separate package). The fact that *abc* itself didn't have to be changed at all makes the extension robust with respect to *abc* upgrades. Also, since the new AST node extends an existing node, very little functionality needs to be reimplemented. The associated interfaces only have to declare the methods specific to the new node's particular functionality.

In the same way, interfaces *PCLocalVars* and *PCCast* were defined, along with implementing classes, for the private pointcut variables and **cast** pointcut extensions. Corresponding production methods were added to the extended AspectJ node factory.

5.4 Adding new frontend passes

Implementing the "global pointcuts" extension described in Section 4.2 requires somewhat more extensive additions to the compiler — first, all global pointcuts need to be collected, and then each pointcut must be replaced with the conjunction of the original pointcut and all applicable global pointcuts.

Polyglot's visitor-based architecture makes implementing that very easy. We add two new passes. The first stores all global pointcuts in a static variable, and the second applies that pointcut to the relevant code. For reasons of code brevity, these two passes are implemented by the same class, *GlobalAspects*; it uses a member variable called *pass* to distinguish which of the two functions it is performing.

The traversal of the AST is performed by the *ContextVisitor* Polyglot class. The new pass extends *ContextVisitor* with a method that performs the required action when it encounters a relevant AST node.

The following code fragment illustrates the behaviour of the new visitor upon entering an AST node:

```
public NodeVisitor enter(Node parent, Node n) {
    if (pass == COLLECT
      && n instanceof GlobalPointcutDecl) {
        ((GlobalPointcutDecl) n).
          registerGlobalPointcut(this, context(), nodeFactory);
    }
    return super.enter(parent, n);
}
```

As mentioned above, both new passes are implemented by the same class, and hence the check that *pass==COLLECT* makes sure we do the right thing. If the current node is a *GlobalPointcutDecl* (one of the new AST nodes defined in section 5.3), we call its special method so it registers itself with the data structure storing global pointcuts. Then we delegate the rest of the work (the actual traversal) to the superclass.

The implementation of the *leave*() method, which is called when the visitor leaves an AST node and has the option of rewriting the

```
public static CastShadowMatch
                matchesAt(MethodPosition pos)
{
    if (!(pos instanceof StmtMethodPosition))
        return null;

    Stmt stmt = ((StmtMethodPosition) pos).getStmt();

    if (!(stmt instanceof AssignStmt))
        return null;
    Value rhs = ((AssignStmt) stmt).getRightOp();

    if (!(rhs instanceof CastExpr))
        return null;
    Type cast_to = ((CastExpr) rhs).getCastType();

    return new CastShadowMatch(
                pos.getContainer(), stmt, cast_to);
}
```

Figure 3: The CastShadowMatch.matchesAt(...) method

node if necessary, is very similar. If *pass==CONJOIN* and we are at an appropriate node, we return the conjunction of the node and the global pointcut.

The sequence of passes that the compiler goes through is specified in the special singleton *ExtensionInfo* class. By subclassing it and inserting our new passes in an overridden method which then calls the original method, we make sure the original sequence of passes is undisturbed. Note that this mechanism makes the extension robust with respect to changes in the base *abc* passes — we can add and rearrange passes without breaking the extension.

5.5 Adding new join points

To implement the cast and throw pointcuts, we first need to extend the list of join point types. This is done by adding to a list of factory objects which the pointcut matcher iterates over to find all join point shadows. The *listShadowTypes* method is defined in the *AbcExtension* class and is overridden for *eaj*:

```
protected List /*<ShadowType>*/ listShadowTypes()
{
    List /*<ShadowType>*/ shadowTypes =
            super.listShadowTypes();
    shadowTypes.add(CastShadowMatch.shadowType());
    shadowTypes.add(ThrowShadowMatch.shadowType());
    return shadowTypes;
}
```

The definitions of *CastShadowMatch* and *ThrowShadowMatch* are very similar and we therefore limit ourselves to discussing the former.

The *CastShadowMatch.shadowType*() method just returns an anonymous factory object which delegates the work of finding a join point to a static method in the *CastShadowMatch* class. This method, *matchesAt*(...), takes a structure describing a position in the program being woven into and returns either a new object representing a join point shadow or *null*; the code for it is given in Figure 3.

The purpose of the *MethodPosition* parameter is to allow *abc* to iterate through all the parts of a method where a join point shadow can occur, and ask each factory object whether one actually does. There are four types of *MethodPosition* for normal AspectJ shadows:

- Whole body shadows: execution, initialization, preinitialization
- Single statement shadows: method call, field set, field get
- Statement pair shadows: constructor call
- Exception handler shadows: handler

Most shadows either fall into the category of "whole body" or "single statement". Two are special; a constructor call join point encompasses both the *new* instruction that creates the object as well as the *invokespecial* that initialises it, and handler join points can only be found by looking at the exception handler table for a method, rather than its statements.

If a new join point requires an entirely new kind of method position, then the code that iterates over them can be overridden.

The first job of the *matchesAt*(...) method is to check that we are at the appropriate position for a **cast** pointcut, namely one with a single statement. Next, we need to check whether there is actually a cast taking place at this position; the grammar of Jimple makes this straightforward, as a cast operation can only take place on the right-hand side of an assigment statement. If no such operation is found, we return *null*; otherwise we construct an appropriate object.

Defining the *CastShadowMatch* class also requires a few other methods, connected with defining the correct values to be bound by an associated **args** pointcut, reporting the information required to construct a *JoinPoint.StaticPart* object at runtime, and recording the information that a pointcut matches at this shadow in an appropriate place for the weaver itself to use. The details are straightforward, and we omit them for reasons of space.

5.6 Extending the pointcut matcher

Again, we describe the implementation of the **cast** pointcut and omit discussion of the almost identical throw pointcut. Once the corresponding join point shadow has been defined, writing the appropriate backend class is straightforward. The pointcut matcher tries every pointcut at every join point shadow found, so all the **cast** pointcut has to do is to check whether the current shadow is a *CastShadowMatch*, and if so verify that the type being cast to matches the *TypePattern* given as argument to the **cast** pointcut:

```
protected Residue matchesAt(ShadowMatch sm)
{
    if (!(sm instanceof CastShadowMatch))
        return null;
    Type cast_to = ((CastShadowMatch) sm).getCastType();

    if (!getPattern().matchesType(cast_to))
        return null;

    return AlwaysMatch.v();
}
```

The *AlwaysMatch.v*() value is a *dynamic residue* that indicates that the pointcut matches unconditionally at this join point. For those pointcuts where matching cannot be statically determined, this is replaced by one which inserts some code at the shadow to check the condition at runtime.

5.7 Extending the runtime library

AspectJ provides dynamic and static information about the current join point through *thisJoinPoint* and associated special variables.

For the **cast** pointcut extension, this runtime interface was extended to reveal the signature of the matching cast. For example, the following aspect picks out all casts (except for the one in the body of the advice) and uses runtime reflection to display the type that is being cast to at each join point:

```
import org.aspectbench.eaj.lang.reflect.CastSignature;

aspect FindCasts
{
    before():
```

```
        cast(*) && !within(FindCasts)
    {
        CastSignature s = (CastSignature)
            thisJoinPointStaticPart.getSignature();

        System.out.println("Cast to: " +
                s.getCastType().getName());
    }
}
```

Implementing this requires changes both in the backend of the compiler (where the static join point information is encoded for the runtime library to read later), and the addition of new runtime classes and an interface.

Static join point information is encoded in a string which is parsed at runtime by a factory class to construct the objects accessible from *thisJoinPointStaticPart*. This happens just once, namely in the static initialiser of the class where the join point shadow is located. The alternative, which is to directly generate code to construct these objects, would be expensive in terms of the size of the bytecode produced; using strings provides a compact representation without too much runtime overhead.

The static information for a **cast** pointcut is encoded as follows. To allow us to easily reuse the existing parser for such strings, a fair amount of dummy information is generated, corresponding to properties that cast join points do not have. For example, modifiers such as **public** are important for join points that have a method or field signature associated with them, but make no sense for the cast join point. The string for the **cast** pointcut is constructed from four parts:

- Modifiers (encoded as an integer — 0 for a cast)
- Name (usually a method or field name, but for a cast it is just "cast")
- Declaring type — class in which the join point occurs
- Type of the cast

For example, a cast join point within a method in the class *IntHashTable* which casts the value retrieved from a *HashMap* to an *Integer* would produce the following encoded string:

`"0-cast-IntHashTable-Integer"`

The runtime factory is subclassed to add a method that creates an object implementing the new *CastSignature* interface for appropriate join points. The aforementioned *AbcExtension* class has a method which specifies which runtime class should be used as a factory for **thisJoinPointStaticPart** objects, which is overriden so that runtime objects are created with the new factory:

```
public String runtimeSJPFactoryClass()
{
    return
        "org.aspectbench.eaj.runtime.reflect.EajFactory";
}
```

5.8 Code measurements

To enable the reader to assess the amount of effort involved in implementing each of these new features, we have summarised some statistics in Figure 4. The table shows the size of the whole parser, and of the boilerplate for factories in the top and penultimate row, respectively. The most interesting part is the breakdown by construct in the middle. For private pointcut variables, all the work goes into defining new AST nodes, and there is no need to define new passes or to touch the weaver in any way. By contrast, global pointcuts require the introduction of new Polyglot passes, which reduce the new construct to existing AspectJ constructs. Finally, for cast and throw pointcuts, there is substantial work in the weaver, because these introduce a new type of join point.

***eaj* measurements**		Files	Lines of code
Parsing		1	74
Private pointcut variables	AST nodes	2	130
	Passes	0	0
	Weaver	0	0
	Runtime	0	0
Global pointcut declarations	AST nodes	4	64
	Passes	1	77
	Weaver	0	0
	Runtime	0	0
Cast poincuts	AST nodes	2	46
	Passes	0	0
	Weaver	2	94
	Runtime	2	27
Throw pointcuts	AST nodes	2	46
	Passes	0	0
	Weaver	2	91
	Runtime	2	16
Extension information and shared classes		7	205
Total		27	870

Figure 4: Code measurements for *eaj*

Throw-pointcut statistics	*ajc*	*abc*
Core compiler/runtime files modified	8	0
throw-specific files created	2	6
Extension-specific factories modified	-	5
Total files touched	10	11
Lines of code written[1]	103	187

Figure 5: The throw pointcut in *ajc* and *abc*.

It is pleasing to us that the distinction between the examples is so sharp, as it gives good evidence that the aim of modularity has been achieved. We believe that the amount of code that needs to be written also meets the criterion of proportionality that was introduced at the beginning of this paper. The criterion of simplicity is more difficult to measure, but we hope that the sample code in this section suffices to convince the reader that we have succeeded in this respect as well. The examples presented here do not demonstrate analysis capability: assessment of that criterion is ongoing work.

6. RELATED WORK

Others before us have identified the need for a workbench to support the rapid developments in aspect-oriented programming language research. In this section, we review a number of such proposals, and contrast them with the approach taken in *abc*.

ajc. The *de facto* standard workbench for research into variations and extensions of AspectJ is the *ajc* compiler. It has served this purpose admirably well, and for example [21,24] report on the successful integration of substantial new features into *ajc*.

We believe that, in view of the explosion of research into new features and analyses, the time has now come to disentangle the

[1]Note that the numbers in Figure 5 for *abc* take into account the relevant lines of files which are listed under "Extension information and shared classes" in Figure 4.

code of the base compiler from that of the extensions. The benefits are illustrated by the table in Figure 5. It compares the implementation of the **throw** pointcut in *abc* and *ajc*. In the case of *ajc*, we have to modify a large number of existing files, thus tangling the new extension with the existing compiler base. At the cost of some factory classes (and thus some more lines of code), *abc* disentangles the two completely.

A more detailed comparison of the characteristics of *ajc* and *abc* can be found in a companion paper [3]; in particular, that paper explores the fact that *abc*'s compile times are typically a factor 3-4 slower than those of *ajc*.

Javassist. Javassist is a reflection-based toolkit for developing Java bytecode translators [7]. Compared to other libraries such as BCEL, it has the distinguishing feature that transformations can be described using a source-level vocabulary. Compared to *abc*, it provides some of the combined functionality of the Java-to-Jimple translator plus the advice weaver, but its intended applications are different: in particular, it is intended for use at load-time. Consequently, Javassist does not provide an analysis framework like Soot does in *abc*. In principle, such a framework could be added, but it would require the design of a suitable intermediate representation akin to Jimple.

Josh. Josh is an open implementation of an AspectJ-like language based on Javassist [6], and as such it is much closer in spirit to *abc*. Indeed, the primary purpose of Josh is to experiment with new pointcut designators, although it can also be used for features such as parametric introductions. Because of the implementation technology, there is no special support for the usual static checks in the frontend, which is provided in *abc* by the infrastructure of Polyglot. Josh does not cover the whole of AspectJ, which limits its utility in realistic experiments.

Logic meta-programming. A more radical departure from traditional compiler technology is presented by *logic meta-programming*, as proposed by [9, 14]. Here, program statements where extra code should be woven in are selected by means of full-fledged Prolog programs. This adds significant expressive power, and like Josh, the design makes it easy to experiment with new kinds of pointcuts. The system operates on abstract syntax trees, which are not a convenient representation for transformation and analysis — many years of research in the compilers community have amply demonstrated the merits of a good intermediate representation. A further disadvantage, in our view, is the lack of static checks due to the increased expressive power. The success of AspectJ can partly be explained by the fact that it provides a *highly disciplined* form of meta-programming; some of that discipline is lost in logic meta-programming, because the full power of Prolog precludes certain static checks. Nevertheless, a system based on these ideas is publicly available [27], and it is used as a common platform by a number of researchers.

Pointcuts as functional queries. Eichberg, Mezini and Ostermann have very recently suggested an open implementation of pointcuts, to enable easy experimentation with new forms of pointcuts [12]. Their idea is closely related to that of logic meta-programming, namely to use a declarative query language to identify join point shadows of interest. A difference is that they opt for the use of the XML query language XQuery instead of a logic language. Furthermore, [12] only deals with static join points. As argued in the introduction, several recent proposals for new pointcut primitives require data flow analyses. We believe that it is not convenient to express such analyses via queries on syntax trees. It is however quite easy to transfer some of the ideas of [12] to *abc*, by letting the queries range over Polyglot ASTs. A challenge, then, is to define appropriate type rules to implement as part of the frontend.

7. CONCLUSIONS AND FUTURE WORK

We have presented *abc*, and its use as a workbench for experimentation with extensions of AspectJ. Our primary design goal was to completely disentangle new features from the existing codebase, and this goal has been met. We hope that such disentangling will enable yet more rapid developments in the design of aspect-oriented programming languages, and the integration of ideas from multiple research teams into a single system, where the base can evolve independently of the extensions.

This project has also been an evaluation of the extensibility of Polyglot and Soot, from the perspective of aspect-oriented software development. We now summarise their role in the extensibility of our design, and identify possible improvements.

Polyglot. Polyglot turned out to be highly suited to our purposes. Its extension mechanisms are exactly what is needed to implement AspectJ itself as an extension of Java, with only minimal code duplication. This in turn makes the development of *abc* relatively independent of further improvements to Polyglot.

As we have remarked earlier, the Polyglot mechanism of *delegates* mimicks that of ordinary intertype declarations, whereas *extension nodes* roughly correspond to what an AspectJ programmer would naturally do via **declare parents** and interface intertype declarations. Polyglot achieves this effect by cunningly creating a replica of the inheritance hierarchy in code, which then provides the hooks for appropriate changes. Arguably that mechanism is somewhat brittle, and it is certainly verbose, replicating the same information in multiple places of the code.

We thus face the question whether it would be possible to extend *abc* using AspectJ, or indeed any other dialect of Java that features open classes. The answer is in the positive, as *abc* is written in pure Java. It follows that users who prefer to use AspectJ to extend *abc* can do so without further ado.

Would the result be more compact and understandable code? Unfortunately, a significant proportion of Polyglot's extensions is taken up by boilerplate code for generic visitors in each new AST node. To generate that automatically, one would need reflection or a feature akin to parametric introductions [15]. The reflection route has been used with much success, in a framework by Hanson and Proebsting [16] that is very similar to Polyglot.

Soot. The choice of Soot as the basis for our code generation and weaver has had a fundamental impact not only on the quality of the code that is generated, but also on the ease by which the transformations are implemented. The Jimple intermediate representation of Soot has been honed on a great variety of optimisations and analyses before we applied it to *abc*, and we reap the benefits of this large body of previous work.

Equally important has been the use of the Dava decompiler that is part of the Soot framework. This makes it much easier to pinpoint potential problems, and to communicate the ideas about code generation to others. It also opens the way to exciting new visualisations, for example to indicate at source level exactly what dynamic residue was inserted at a join point shadow.

In this paper, we have not yet fully exploited the analysis capabilities of Soot. In particular, for the optimised implementation of advanced features such as predicted control flow [19], data flow pointcuts [21] and trace cuts [10, 28], the interprocedural analyses of Soot will be very important. In such cases, one needs to first weave naively, in order to get an approximation of the relevant analysis structures. These are then used to weave again, but now more precisely, removing redundant residues. We report on the details of this process (which also applies to the implementation of **cflow** [25]) in a companion paper [4].

The ultimate test of *abc*'s extensibility will be its use by others. We are encouraged that at the time of writing, at least 10 academic groups have started to use *abc* for their own research, including the projects described in [17, 20].

8. REFERENCES

[1] abc. The AspectBench Compiler. Home page with downloads, FAQ, documentation, support mailing lists, and bug database. `http://aspectbench.org`.

[2] Jonathan Aldrich. Open modules: A proposal for modular reasoning in aspect-oriented programming. Technical Report CMU-ISRI-04-108, Institute for Software Research, Carnegie Mellon University, 2004.

[3] Pavel Avgustinov, Aske Simon Christensen, Laurie Hendren, Sascha Kuzins, Jennifer Lhoták, Ondrej Lhoták, Oege de Moor, Damien Sereni, Ganesh Sittampalam, and Julian Tibble. Building the abc AspectJ compiler with Polyglot and Soot. Technical Report abc-2004-4, aspectbench.org, 2004.

[4] Pavel Avgustinov, Aske Simon Christensen, Laurie Hendren, Sascha Kuzins, Jennifer Lhoták, Ondrej Lhoták, Oege de Moor, Damien Sereni, Ganesh Sittampalam, and Julian Tibble. Optimising AspectJ. Technical Report abc-2004-3, aspectbench.org, 2004.

[5] Michael Brukman and Andrew C. Myers. PPG: a parser generator for extensible grammars, 2003. Available at `www.cs.cornell.edu/Projects/polyglot/ppg.html`.

[6] Shigeru Chiba and Kiyoshi Nakagawa. Josh: an open AspectJ-like language. In K. Lieberherr, editor, *3rd International Conference on AOSD*, pages 102–111, 2004.

[7] Shigeru Chiba and Muga Nishizawa. An easy-to-use toolkit for efficient Java bytecode translators. In *2nd International coference on Generative Programming and Component Engineering (GPCE '03)*, volume 2830 of *LNCS*, pages 364–376, 2003.

[8] Adrian Colyer and Andrew Clement. Large-scale AOSD for middleware. In *3rd International Conference on AOSD*, pages 56–65. Association for Computing Machinery, 2004.

[9] Kris de Volder. Aspect-oriented logic meta-programming. In Pierre Cointe, editor, *2nd International Conference on Meta-level Architectures and Reflection*, volume 1616 of *LNCS*, pages 250–272, 1999.

[10] Rémi Douence, Pascal Fradet, and Mario Südholt. Composition, reuse and interaction analysis of stateful aspects. In Karl Lieberherr, editor, *3rd International Conference on AOSD*, pages 141–150, 2004.

[11] Chris Dutchyn, Gregor Kiczales, and Hidehiko Masuhara. Tutorial: AOP language exploration using the Aspect Sand Box. In *1st International Conference on AOSD*, 2002.

[12] Michael Eichberg, Mira Mezini, and Klaus Ostermann. Pointcuts as functional queries. In *Second ASIAN Symposium on Programming Languages and Systems (APLAS 2004)*, Springer Lecture Notes in Computer Science, 2004.

[13] Etienne Gagnon, Laurie J. Hendren, and Guillaume Marceau. Efficient inference of static types for Java bytecode. In *Static Analysis Symposium*, pages 199–219, 2000.

[14] Kris Gybels and Johan Brichau. Arranging language features for more robust pattern-based crosscuts. In *2nd International Conference on AOSD*, pages 60–69. ACM Press, 2003.

[15] Stefan Hanenberg and Rainer Unland. Parametric introductions. In Mehmet Aksit, editor, *2nd International Conference on AOSD*, pages 80–89, 2003.

[16] David Hanson and Todd Proebsting. A research C# compiler. *Software — Practice and Experience*, to appear, 2004.

[17] Bruno Harbulot and John R. Gurd. Using AspectJ to separate concerns in parallel scientific Java code. In *Proceedings of the 3rd international conference on AOSD*, pages 122–131. ACM Press, 2004.

[18] Jim Hugunin. Guide for developers of the AspectJ compiler and weaver, 2004. Available at `http://dev.eclipse.org/viewcvs/index.cgi/~checkout~/org.aspectj/modules/docs/developer/compiler-weaver/index.html?rev=1.1&content-type=text/html&cvsroot=Technology_Project`.

[19] Gregor Kiczales. The fun has just begun. Keynote address at AOSD. Available at `aosd.net/archive/2003/kiczales-aosd-2003.ppt`, 2003.

[20] Shriram Krishnamurthi, Kathi Fisler, and Michael Greenberg. Verifying aspect advice modularly. In *ACM SIGSOFT International Symposium on the Foundations of Software Engineering*, 2004.

[21] Hidehiko Masuhara and Kazunori Kawauchi. Dataflow pointcut in aspect-oriented programming. In *1st Asian Symposium on Programming Languages and Systems*, volume 2895 of *LNCS*, pages 105–121, 2003.

[22] Hidehiko Masuhara, Gregor Kiczales, and Chris Dutchyn. A compilation and optimization model for aspect-oriented programs. In *Compiler Construction*, volume 2622 of *LNCS*, pages 46–60, 2003.

[23] Nathaniel Nystrom, Michael R. Clarkson, and Andrew C. Myers. Polyglot: An extensible compiler framework for Java. In *12th International Conference on Compiler Construction*, volume 2622 of *LNCS*, pages 138–152, 2003.

[24] Kouhei Sakurai, Hidehiko Masuhara, Naoyasu Ubayashi, Saeko Matsuura, and Seiichi Komiya. Association aspects. In Karl Lieberherr, editor, *3rd International Conference on AOSD*, pages 16–25, 2004.

[25] Damien Sereni and Oege de Moor. Static analysis of aspects. In *Proceedings of the 2nd International Conference on AOSD*, pages 30–39, 2003.

[26] Raja Vallée-Rai, Etienne Gagnon, Laurie J. Hendren, Patrick Lam, Patrice Pominville, and Vijay Sundaresan. Optimizing Java bytecode using the Soot framework: Is it feasible? In *Compiler Construction, 9th International Conference (CC 2000)*, pages 18–34, 2000.

[27] Kris De Volder. The TyRuBa metaprogramming system. Available at `http://tyruba.sourceforge.net/`.

[28] Robert Walker and Kevin Viggers. Implementing protocols via declarative event patterns. In *ACM Sigsoft International Symposium on Foundations of Software Engineering (FSE-12)*, 2004.