

Beyond annotations : a proposal for extensible java (XJ).

CLARK, Anthony <<http://orcid.org/0000-0003-3167-0739>>, SAMMUT, Paul and WILLANS, James

Available from Sheffield Hallam University Research Archive (SHURA) at:

<http://shura.shu.ac.uk/11902/>

This document is the author deposited version. You are advised to consult the publisher's version if you wish to cite from it.

Published version

CLARK, Anthony, SAMMUT, Paul and WILLANS, James (2008). Beyond annotations : a proposal for extensible java (XJ). In: CORDY, James R. and ZHANG, Lu, (eds.) Proceedings Eighth IEEE International Working Conference on Source Code Analysis and Manipulation. IEEE, 229-238.

Copyright and re-use policy

See <http://shura.shu.ac.uk/information.html>

Middlesex University Research Repository

An open access repository of
Middlesex University research

<http://eprints.mdx.ac.uk>

Clark, Tony and Sammut, Paul and Willans, James (2008) Beyond annotations: a proposal for extensible java (XJ). In: Source Code Analysis and Manipulation, 2008 Eighth IEEE International Working Conference on 28-29 Sept. 2008. IEEE, pp. 228-238.

<http://dx.doi.org/10.1109/SCAM.2008.34>

Available from Middlesex University's Research Repository at
<http://eprints.mdx.ac.uk/5817/>

Copyright:

Middlesex University Research Repository makes the University's research available electronically.

Copyright and moral rights to this thesis/research project are retained by the author and/or other copyright owners. The work is supplied on the understanding that any use for commercial gain is strictly forbidden. A copy may be downloaded for personal, non-commercial, research or study without prior permission and without charge. Any use of the thesis/research project for private study or research must be properly acknowledged with reference to the work's full bibliographic details.

This thesis/research project may not be reproduced in any format or medium, or extensive quotations taken from it, or its content changed in any way, without first obtaining permission in writing from the copyright holder(s).

If you believe that any material held in the repository infringes copyright law, please contact the Repository Team at Middlesex University via the following email address:

eprints@mdx.ac.uk

The item will be removed from the repository while any claim is being investigated.

Beyond Annotations: A Proposal for Extensible Java (XJ)

Tony Clark, Paul Sammut, James Willans - Ceteva Ltd.

1 Introduction

In his 1994 paper on Language Oriented Programming (LOP) [1], Martin Ward proposes that the problems of complexity, conformity, change and invisibility that occur with large software systems can be addressed by designing a *formally specified, domain-oriented, very high level programming language* as the basis for system design. Recently, Guy Steele [2] has emphasised this point:

a good programmer in these times does not just write programs. [...] a good programmer does language design, though not from scratch, but building on the frame of a base language

The term Domain Specific Language (DSL) has been used to refer to languages that have been designed for a restricted class of applications. Martin Fowler [3] makes the distinction between *external* DSLs and *internal* DSLs. An external DSL is written outside the main language of an application whereas an internal DSL is written in and uses the main application language.

Medium to large scale enterprise information systems are implemented using a collection of different technologies. These include JavaScript, JUnit, Ant, (lots of) XML, Drools and various extensions to Java implemented using annotations. In addition to fairly standard technologies such as those mentioned above, there are a large number of extensions being proposed to standard programming languages such as Java. Many of these extensions are implemented using different pre-processor technologies. Each of these different technologies addresses a different aspect of the overall system. Therefore the system is implemented using a collection of loosely integrated DSLs.

The arguments made by Ward in favour of LOP and DSLs are even more relevant today given the proliferation of technologies and the scale of software systems. Developers are using these techniques all the time in various different ways.

The danger with home-grown LOP technology is that it can make the problem worse, not better. If developers use a variety of methods and tools to construct a language then they may get some immediate benefit, but issues will arise if the underlying technology




Fig. 1: Abstract Model of Development

is not stable. In addition, to be useful DSLs must be portable - perhaps not as widely distributed as general purpose languages, but transferrable nonetheless. For large scale benefit, it should not be necessary to supply the complete language support system along with a system module.

This paper is a proposal for a language extension to Java called XJ that supports LOP and therefore allows DSLs to be constructed to be standard, portable and easily understood. The extension is conservative in the sense that it will not conflict with any existing language features and will preserve backward compatibility. The proposal extends classes with syntax definitions to produce new modular language constructs, called syntax-classes, that can easily be distributed along with an application in the usual way.

The LOP extensions to Java have been implemented in an existing language called XMF [4]. XMF is a high-level object-oriented language designed to support DSLs. XMF has been used in a commercial product and has a Java-like sub-language. Therefore we are confident that the XJ proposal has been adequately validated.

This paper is structured as follows: section 2 reviews the state of the art in DSLs; section 3 describes the XJ language extensions; section 4 describes some LOP defined examples using XJ; finally, section 5 analyses the proposal and compares it to related proposals for Java-based LOP technologies.

2 Domain Specific Languages

Consider an idealised development process. You are given an application to develop. You start by trying to understand the kind of executions you will need to perform - just fragments at first. Gradually, you join

the executions together to form ever greater descriptions of what will need to happen. The executions, are described in terms of entities from the application domain: financial transactions, customers, interest rates and so on.

You start to notice common themes in the executions and design execution rules that cover them. Gradually, you wind up with a fairly comprehensive collection of execution rules; each rule expects to be given some information that causes it to perform some calculation. The rules define an idealised *domain specific execution engine* and the data that drives the engine defines a *domain specific program*. The range of variability in the programs you can given the engine is a *domain specific language*.

Once you have arrived at an engine and language then you choose a *general purpose language*. You understand the language very well and your implementation task becomes one of working out how to translate the domain specific program into a real-world program so that it faithfully executes on your *general purpose engine*. See figure 1.

Of course, the above process is idealised, however it is representative of how a developer gets from an application specification to a working implementation. How closely the process is followed will depend on how expert the programmer is (there are lots of translations going on). An expert programmer will perform many of the steps instinctively and perhaps often not even be aware of the various representations.

The difference between the DSL and the implementation language is the *representation chasm*. Once the process has jumped from one side of the chasm to the other, it is very difficult to get back. This one-way process is a source of many problems since it loses the meaning of the original representation. Once lost, it is difficult for humans to interpret the code (maintenance becomes an issue), and it is difficult for machines to extract meaning from the code (for example building reuse libraries or generating efficient code).

DSLs have their roots in the language extensions mechanisms provided by Lisp and Scheme. These languages promoted a declarative approach to application development through the use of macros. A key feature for LOP is the ability to process the abstract syntax of the underlying language. The Lisp family of languages represents both data and programs using S-expressions which has a simple structure.

Until the development of C++ templates, most mainstream languages other than Lisp, provided very little support for LOP. The C pre-processor is not really a suitable technology for LOP. Whilst C++ templates provide much more support for LOP than the C pre-processor, they do not support access to the underlying abstract syntax.

Current mainstream languages do not provide any

support for LOP, however there is increasing interest in LOP and DSLs because of the perceived benefits outlined above. A number of language extensions to Java have been proposed to support LOP and these are analysed in section 5.

Many of the approaches to LOP involve language pre-processors that perform source-to-source translations. Whilst this provides LOP and DSLs, it introduces a new technology, usually a parser and some form of source code translation mechanism, that is not integrated with the main implementation language. In order to use the new language in Java, the provider must supply the pre-processing technology which may not work with Java development tools and will be dependent on a third party for maintenance.

The structure of this section is as follows: section 2.1 describes a number of architectural styles for developing DSLs; section 2.2 describes a number of technologies used to implement DSLs; section 2.3 analyses the styles and technologies and makes a proposal for a standard extension to Java that supports LOP and DSLs; finally, section 2.4 lists a number of example DSLs that could be added to Java using the proposals.

2.1 DSL Architectural Styles

There are a number of approaches to architecting a DSL using LOP technologies. This section briefly describes the main styles.

2.1.1 Internal vs External DSLs

Fowler makes the distinction between *internal* and *external* DSLs. An internal DSL makes use of an existing language and possibly integrates with it whereas an external DSL is completely separate. It should be noted that, if technology supports a fully integrated internal approach then that can be used to implement an external approach. This is because such a technology would allow the complete replacement of the host language syntax whilst still using the host language for the implementation of the DSL.

2.1.2 DSL Integration

DSLs may be *fully* or *partially integrated* with their host language. A full integration allows the DSL constructs to be interleaved, where appropriate, with the host constructs. Effectively the host language is extended by the DSL. Partial integration involves levels of restriction on the potential for interleaving. A very restricted integration allows the DSL constructs to occur at pre-defined points in the host language. A fully restricted integration is effectively an external DSL.

2.1.3 Concrete Syntax

DSLs may introduce *new syntax constructs*, may re-interpret the syntax constructs of the host language, or may simply use the existing constructs of a language in a systematic way. New syntax constructs are perhaps the purest form of DSL but will require some parsing technology. Re-interpretation often involves observing patterns of existing language constructs and treating them specially in some way, perhaps involving pattern matching technology. A typical example is to treat certain procedure calls, method calls or field references in a special way.

2.1.4 Translational vs Execution

DSLs may be *translational* or may involve their own *execution engines*. A translational DSL transforms source code from the DSL constructs to existing constructs in a target language. This may involve a pre-processor or be integrated within the target language compiler. A translational approach requires technology for manipulating syntax - concrete, abstract or both.

DSLs that have their own execution engines tend to be external DSLs (but need not be). In this case the DSL program is a data structure that is processed by an engine written specifically for the purpose. Such an approach is related to software frameworks as discussed below.

2.2 Existing Technologies for DSLs

This section describes a number of different technologies and approaches that are currently used to implement DSLs.

2.2.1 XML

XML has been proposed as a technology for expressing DSLs. This is possible and achieves a standard representation, however XML is not integrated with a programming language in any way and cannot take advantage of the feature of the language in the DSL constructs, therefore it does not meet the criteria set out in [2] for language definition. In certain limited highly declarative domains (for example expressing data structure relationships, declaring project dependencies or menu items) XML is fine, but otherwise it requires a complete language engine to be written from scratch using a representation that is not particularly human-friendly.

2.2.2 Pre-processors

A pre-processor can be used to add DSL features to an existing language. Some pre-processors also allow macros to be defined so that the user can define

new language features. In both cases this achieves the aims of LOP but introduces new technology. The pre-processor technology must be distributed with the language definition in order for anyone to use it. Most pre-processing technologies have restricted knowledge of the target language and cannot provide new language features with information about their static context (variables in scope, types etc). This makes it difficult to use pre-processors to extend an existing language in order to achieve [2].

Pre-processors may also be limited to the existing syntax of the target language, i.e. they can interpret patterns of existing constructs, but not introduce new constructs. This feature limits the ability of pre-processors to achieve the goals set out in [1]. Finally, pre-processors may not work with existing language tools. For example, Java IDEs can only process standard Java; language constructs introduced via a pre-processor will not be legal in such an IDE and the code produced by the pre-processor will not be understood by the developer.

2.2.3 Chained Calls

A good place to start when defining a DSL is to make a list of the nouns and verbs associated with the domain. These are good candidates for the data and language features that the DSL will need to support. A style of OOP uses chains of method calls to implement these features. Proponents of this style, claim that the approach achieves the benefits of a DSL without needing any new language features.

As an example of chained calls in Java consider the following example of XML (taken from <http://www.infoq.com/articles/internal-dsls-java>):

```
String sql =
    "select id, name " +
    "from customers c, order o " +
    "where " +
    "c.since >= sysdate - 30 and " +
    "sum(o.total) > " + significantTotal + " and " +
    "c.id = o.customer_id and " +
    "nvl(c.status, 'DROPPED') != 'DROPPED'";
```

This can be refactored as a chained method call DSL:

```
Table c = CUSTOMER.alias();
Table o = ORDER.alias();
Clause recent =
    c.SINCE.laterThan(daysEarlier(30));
Clause hasOrders =
    o.TOTAT.sum().isAbove(significantTotal);
Clause ordersMatch =
    c.ID.matches(o.CUSTOMER_ID);
Clause activeCustomer =
    c.STATUS.isNotNullOr("DROPPED");
```

```
String sql =
    CUSTOMERS
        .where(recent.and(hasOrders)
            .and(ordersMatch)
            .and(activeCustomer)
            .select(c.ID, c.NAME)
            .sql());
```

Whilst chained-calls can lead to readable code and no doubt the approach can be good practice since it focusses design attention on the domain, there is limited scope for proper encapsulation of new concepts and no scope for conveying semantic intent to the language tools.

2.2.4 Processing Concrete Syntax

Systems that support an LOP approach often provide mechanisms for working with concrete syntax. These include patterns for matching structure and templates for generating code. It is important that processing is not limited to concrete syntax since there are often constructs in a language that have no corresponding concrete syntax language feature. For example, in Java there is no concrete syntax for defining a package containing classes. However, in order to achieve [1] it is necessary to be able to produce new features that abstract commonly occurring language patterns.

Systems that are limited to just pattern matching over concrete syntax cannot introduce new language features whose structure differs from that of existing features. In order to deal with new concrete syntax, some form of parsing mechanism is required. One way of achieving this is to allow grammars to be defined within the language. Each new grammar corresponds to a new language construct. In order for the new language constructs to be embedded within the existing language, the grammar rules should be able to extend those of existing constructs.

2.2.5 Scripting and External DSLs

An external DSL is one that does not extend or use the main implementation language. Such a language is often used to implement a module in a larger system. There are a number of scripting languages, such as JRuby, that can be used in conjunction with general purpose programming languages, such as Java. In these cases the GPL is used to implement the core system and the scripting languages provides a DSL for specific aspects, in the case of JRuby this could be the user interface.

For specific types of application this can work well, but it does not really address the key motivations behind [2, 1]. The scripting language is often just another GPL that happens to be used for a specific aspect of the system. The main GPL (Java) knows little about the

semantics of the constructs it is calling in the scripting language. In many cases the scripting language is quite broad in scope and therefore is not really achieving the aim of providing a DSL.

2.2.6 Java Annotations

Java has recently been extended with annotations which can be used to introduce properties to specific places in the source code. The motivation for the annotations has been to provide a mechanism for extending the Java language with new constructs, by allowing existing constructs to have properties. A new class of construct can be defined by distinguishing between existing constructs based on their property annotations. In addition, the annotations are a standard part of the language and therefore can be processed by Java language tools such as compilers and IDEs.

This approach has been used in a number of extensions to Java and, if used lightly, can be successful. However the aim of achieving new program constructs that reflect the domain or capture programming idioms is not achieved using annotations since they cannot introduce any new syntax. In addition, if annotations are overused then the source code becomes difficult to read.

2.2.7 Frameworks and Libraries

A Software Framework is a collection of classes that capture a design pattern or an aspect of an application. The key feature of a software framework is that it represents a reusable component where some of the features of the framework are fixed in structure and behaviour and some are left open to extension by the user. A framework differs from a more traditional software library in that the user provides the framework with the implementation of services that the framework calls. In a software library this works the other way round where the user calls services provided by the library (the so-called *Inversion of Control*). The extension points in a framework are typically supplied with objects whose classes implement abstract methods. At appropriate points in its life-cycle, the framework will invoke the methods of the registered objects.

A DSL is related to frameworks in a number of ways. A DSL that extends the syntax of a language typically uses a parsing framework to deal with the new concrete syntax. If the new syntax is fully integrated with an existing language then extension points can be provided by the parser for the existing language to support the extensions. If the compiler or interpreter for an existing language is a framework then it can provide extension points for new language constructs to be expanded into existing language constructs. Finally, if a DSL is to be implemented as an external language then the execution engine for that language can be viewed as a

framework where the extension points are implemented in data that is supplied as the DSL program.

2.3 Analysis

We have described DSLs, their architectural styles and some technologies that can be used to implement them. This section reviews the benefits and drawbacks of DSLs and then makes some recommendations regarding an approach to DSL architecture and implementation technology. XJ supports all of these recommendations.

2.3.1 Benefits of DSLs

LOP is important because it allows languages to grow in line with [2]. It allows new programming idioms and patterns to be captured in a standard way, can significantly reduce the number of technologies required to produce a large application, and it supports DSL development. The amount of investment necessary to implement a new language from scratch or even to make significant modifications to an existing language is huge. Given that LOP and DSLs are generally accepted to have benefits, this raises a significant problem regarding DSL development.

DSLs improve the readability and maintainability of code since they close the representation chasm. In addition DSLs improve the scope for a developer to convey the intent of program code to system tools such as IDEs and compilers. This is important for tool support of DSLs and to ensure efficient execution of DSL programs. DSLs can be designed so that system component properties are made explicit which provides scope for reuse.

2.3.2 Drawbacks of DSLs

Current approaches to DSLs suffer from having no standard technology that supports language definition. DSL developers are forced to use a variety of technologies or write their own. This has led to the objection that DSLs are high-risk since they require specialised skills for development and are difficult to maintain. Lack of standards for DSLs also mean that tool support for development is weak.

2.3.3 Recommendations

Standardisation is key to achieving the aims of LOP and DSLs. Without standardisation, each LOP developer uses different approaches and technologies which makes the development highly specialized and the maintenance difficult to control. If LOP technology can be standardized then languages become easier to

develop since tools can support the standard and problems of maintenance are greatly reduced. Standardization can also significantly reduce the number of technologies necessary to support LOP and develop DSLs.

Should such a standard support fully integrated DSLs or external DSLs? Given that technology for fully integrated DSLs subsumes that of external, the standard should support a the definition of language constructs that can be fully integrated with a host language. In addition, the standard should allow full concrete syntax extension and the ability to process abstract syntax in arbitrary ways since this is perhaps the essence of DSLs.

Should the standard support a translational approach or an execution engine approach? As noted above, the execution engine approach is universal, but quite specialized. It is more suited to external DSLs than internal DSLs. Whilst there are benefits from an execution engine approach (notably control and debugging), a translational approach is more accessible and supports a more lightweight approach for non-language specialists.

Our proposal is to incorporate the above recommendations as a standard language extension to Java. The extension produces a language called eXtensible Java (XJ) which is introduced in the rest of this paper.

2.4 Example DSLs

The following is a list of some of the DSLs that could be supported by XJ. Many of these examples are existing DSLs that are implemented using the techniques described above.

- | | |
|----------------|---|
| SQL | Traditional methods of embedding SQL in Java have used strings. The Java cannot check the SQL and embedding Java variable values within SQL statements is clumsy. A DSL for SQL would embed the language within Java and allow Java expressions to be references within SQL, see [5]. |
| Beans | A single conceptual entity that is implemented in terms of several Java constructs. A DSL for Java Beans would abstract the different constructs into a single language feature. |
| Comprehensions | Java does not provide declarative program support for processing collections. Comprehensions are a standard way of flexibly processing collections and could easily be added as a DSL. |
| Patterns | Java does not support pattern matching. There are various ways of implementing |

patterns, for example by overloading methods with argument patterns. DSLs would allow a new form of method, calling would dispatch to the method whose argument patterns matched the supplied values.

Testing There are a number of proposals for languages to support testing including JUnit and jMock [6]. DSLs allow the details of testing to be hidden and for tests to be expressed declaratively. A DSL for testing could be defined so that tools supported a test driven approach to development.

XML Both XML and HTML can be embedded as concrete languages within Java via a DSL approach leading to a similar integration as provided by JSP.

Architecture Marcus Voelter [7] shows how system architecture can be expressed using a DSL and describes how code can be generated from the DSL. This is representative of a class of DSLs that can be added to Java to express architecture.

Closures There is current discussion regarding how to add closures to Java. DSLs can support the definition of closures.

SPLA Feature models and product line architectures allow single systems to describe a number of products. DSLs can be used to express the systems and their variation points.

Hardware Many hardware systems are controlled by small scripting languages. DSLs would allow these languages to be embedded within Java and to allow the scripts to take advantage of the features of Java.

Rules Rules systems such as Drools [8] already class themselves as DSLs. These rules systems are very useful for encoding business logic which monitor multiple data sources. A standard mechanism for expressing DSLs would provide scope for including a rules system in Java.

GUI Libraries for constructing GUIs tend to require a large amount of code to create, link and configure the various graphical components. DSLs would allow the detail of this code to be hidden and constructed in a standard way.

3 XJ

XJ is a proposed extension to Java that supports LOP. It introduces the idea of a *syntax-class* into Java. A *syntax-class* is a normal Java class that defines a language grammar. When the Java parser encounters an occurrence of a language feature delimited by a *syntax-class*, the class's grammar is used to process the input. If the parse succeeds then the grammar synthesizes an Java abstract syntax tree (AST). An object of type AST has a standard interface that is used by the Java compiler when it processes the syntax. New types of AST can be defined providing that they implement the appropriate interface.

The rest of this section introduces XJ using examples and analyses the main features: section 3.1 specifies a simple language construct; section 3.2 shows how Java is extended with concrete syntax for the new construct; section 3.3 describes the features in XJ for specifying concrete syntax; section 3.4 describes how the new construct is translated into standard Java; section 3.5 describes how abstract syntax is manipulated in XJ; section 3.6 describes some further simple language examples in XJ; finally, 3.7 reviews the main features of XJ.

3.1 The Select Language Construct

Consider a simple language construct in Java that selects an element from a collection based on some predicate. An example of the new construct is shown below:

```
import language mylang.Select;

public Person getChild(Vector<Person> people) {
    @Select Person p from people where p.age < 18 {
        return p;
    } else { return null; }
}
```

The new language construct is called *Select*. A LOP-defined construct in XJ is used by prefixing a reference to the *syntax-class* with the *@*-character. The value *p* is selected from the vector *people* providing that the age of the person is less than 18. If a value can be selected then it is returned otherwise null is returned.

The use of *Select* is equivalent to the following definition:

```
public Person getChild(Vector<Person> people) {
    for(int i = 0; i < people.size(); i++) {
        Person p = people.elementAt(i);
        if(p.age < 18)
            return p;
    }
    return null;
}
```



```
package mylang;

import language java.syntax.Grammar;

import java.syntax.AST;
import java.syntax.Block;
import java.syntax.Context;
import java.syntax.Statement;
import java.syntax.Sugar;
import java.syntax.Type;
import java.syntax.Var;

public class Select extends Sugar {

    private Type type;
    private Var var;
    private AST collection;
    private AST test;
    private Block body;
    private Block otherwise;

    public Select(Type T,String n,AST c,AST t,Block b,Block o) {
        type = T;
        var = new Var(n);
        collection = c;
        test = t;
        body = b;
        otherwise = o;
    }

    @Grammar extends Statement {
        Select ::=
            T = Type
            n = Name
            'from' c = Exp
            'when' t = Exp
            b = Block
            o = ('else' Block | { return new Block(); })
            { return new Select(T,n,c,t,b,o); }.
    }
}
```

Fig. 2: A Select Command (Part 1 of 2)

3.2 Concrete Syntax for Select

A new language construct is defined in XJ by defining a syntax-class. A syntax-class contains a grammar which is used by the Java parser to process the concrete program syntax and to return an abstract syntax tree (AST). Once a syntax-class has been defined, it can be used in program code by referencing the class after the syntax escape character '@'.

Figure 2 shows the concrete syntax part of the syntax-class for Select. The grammar definition describes how to recognize a select statement. A grammar consists of named parse rules. The grammar for Select extends that for Statement which allows Select to be used wherever a statement is expected and also allows Select to reference the parse rules defined for statements (in this case Type, Exp and Block).

The Select rule specifies that a well-formed statement is a type followed by a name, the keyword `from` followed by an expression, the keyword `when` followed by an expression and then a block which is the body of the select. After the body there may be an optional `else` keyword preceding a block.

In each case within the Select rule, the parse elements produce a value that may optionally be associated with names. For example, the type is associated with the name T. In addition, a parse rule can contain Java statements that return a value. These are enclosed in { and }, and may reference any of the names that have been defined to the left of the statement. The final value returned by the Select rule is an instance of the class Select.

3.3 XJ Grammars and Parsing

XJ requires Java classes to be extended with an optional grammar. Classes that contain a grammar definition are referred to as syntax-classes. XJ also requires that the Java parser is extended to allow the current grammar to be changed during a parse.

Referencing a syntax-class after the syntax escape character '@' causes the Java parser to temporarily switch grammars. When the parser encounters `@C ...`, it finds the class C (loading it if necessary), extracts its grammar and continues with the parse using the grammar defined by the class. The starting non-terminal for the class is always the parse rule named C.

The XJ parser maintains a stack of grammars. When an `@C ...` construct is encountered, the grammar for C is pushed onto the stack. If the current grammar succeeds, then the stack is popped, the current value is returned and the parse continues with the grammar at the top of the stack. XJ grammars can be associated with their own tokenizers, but get the standard Java tokenizer by default.

XJ grammars have been implemented and used for a number of years in the XMF system [15]. XMF is

open-source so the parsing algorithms and language for expressing grammars are available for inspection.

3.4 Abstract Syntax for Select

The value synthesized and returned by a grammar must be an instance of `java.syntax.AST`. If the return value is an instance of one of the standard Java AST classes then no special action needs to be taken by the syntax-class. If the return value is an instance of a user-defined syntax-class then that class must implement the AST interface which is used by the compiler to translate the source code into Java VM code. To make this process easier, a user defined syntax-class can extend `java.syntax.Sugar` which implements the AST interface through a method called `desugar`. The `desugar` method is responsible for translating the receiver into an AST for which the interface is defined (typically desugaring into standard Java code).

The syntax-class `Select` extends the class `java.syntax.Sugar` and defines the `desugar` operation in figure 3. The Java compiler needs to process an AST instance in various ways. It achieves this via the AST interface. The class `java.syntax.Sugar` implements the AST interface by calling `desugar` and then performing the appropriate AST operation on the result. The `desugar` operation is supplied with the current compilation context which contains the variables in scope, types etc.

The `Select` syntax-class uses `desugar` to produce the selection code. The particular code depends on the type of the collection so the first thing that `desugar` does it to work out the type of the collection and dispatch to an appropriate desugaring method. Figure 3 shows how vectors are desugared, other types such as arrays are very similar. The `desugarVector` method returns code using quasi-quotes which are explained in the next section.

3.5 Quasi-Quotes

To define the `desugar` method, a syntax-class will generally construct new AST instances. This process is made easy in XJ through the use of quasi-quotes. A quasi-quoted AST is shown below:

```
[| x + <new java.syntax.Int(1)> |]
```

which is equivalent to the Java expression:

```
new java.syntax.BinExp(
  new java.syntax.Var("x"),
  "+",
  new java.syntax.Int(1))
```

The delimiters `[|` and `|]` transform the enclosing concrete syntax into the corresponding abstract syntax.

```

public AST desugar(Context context) {
    Class<T> cType = context.getType(collection);
    if(isVector(cType))
        return desugarVector(cType,context);
    else // More cases...
}

public AST desugarVector(Class<T> cType,Context context) {
    Var done = context.newVar();
    Var coll = context.newVar();
    return [| boolean <done> = false;
        <cType> coll = <collection>;
        for(int i = 0; i < <coll>.size(); i++) {
            <underlyingType(cType)> <var> = <coll>.elementAt(i);
            if(<test>) {
                <done> = true;
                <body>;
            }
        }
        if(!<done>)
            <otherwise>;
    |];
}
}

```

Fig. 3: A Select Command (Part 2 of 2)

Within [| and |] the delimiters < and > can be used to 'unquote' the syntax in order to drop in some abstract syntax. The two forms of delimiters can be arbitrarily nested. Quasi-quotes are an easy way to create code templates in XJ.

The drop-quotes < and > allow a single AST valued expression to be 'dropped' into program code. Sometimes, it is necessary to splice a collection of program elements into a single location. For example, this can occur when a single DSL construct expands into a collection of method definitions. XJ provides splice-quotes <\$ and \$> to support this. For example the following is a method that constructs a class definition with a fixed field and a variable number of methods:

```

public AST mkClass(String n,Vector<Method> M) {}
return
    [| public class <n> {
        private int storage;
        <$ M $>
    }
    |];
}

```

3.6 Other Examples

Syntax expansions can make use of syntax-classes (similar to macros-calling-macros) in XJ. Some of the ex-

amples in the rest of this paper use a couple of syntax-classes in this way. These are specified in this section without their implementations.

3.6.1 Iterate

The UML Object Constraint Language provides a useful language construct called 'iterate' which is a restricted form of collection-folding. The following is a simple example of the use of a syntax-class called Iterate:

```

@Iterate int i in nums with int sum = 0 {
    return sum + i;
}

```

The example adds up a collection of integers in the collection nums. This code is equivalent to a call of a method called addUp:

```

public int addUp(Vector<Integer> nums) {
    int sum = 0;
    for(int i : sums)
        sum = sum + i;
    return sum;
}

```

3.6.2 Comprehensions

A set comprehension is a standard construct in mathematics for processing set-elements and constructing new sets. This is a useful feature for processing Java collections:

```
@Cmp(x + 1) {
    int x <- nums
}
```

which creates a new collection by adding 1 to the integers in the collection `nums`. This is equivalent to a call of `add1`:

```
public Vector<Integer> add1(Vector<Integer> nums) {
    Vector<Integer> vec = new Vector<Integer>();
    for(int x : nums)
        vec.addElement(x + 1);
    return vec;
}
```

3.7 Review

The main features of XJ are:

syntax-classes XJ extends Java classes with the ability to be associated with grammars.

grammars XJ extends Java classes with grammars. A class with a grammar definition can be used as a language construct. Grammars can extend other grammars causing the syntax rules of the parent grammar to be included in the child.

@ A new type of syntax construct is added to Java. Any construct starting with `@` introduces a reference to a syntax-class. The path following the `@` must reference a syntax-class. When the parser encounters an `@` it dispatches to the grammar for the syntax-class.

import A new import mode for importing references to syntax classes. Once imported, a syntax-class can be referenced by name after `@` without further qualification.

AST Access to, and standardization of, Java abstract syntax. The compiler interface for abstract syntax is also standardized along with a compiler context. Various utility classes such as `Sugar` are provided.

quasi-quotes Delimiters `[[` and `]]` that allow abstract syntax to be constructed as though it were concrete syntax. Drop-quotes `<` and `>` within quasi-quotes unquote the delimited expression and allow code templates to be

constructed. The splice-quotes, `<$` and `>` splice sequences of AST objects into a single location within quasi-quotes.

lifting Any Java value can be transformed into an AST object whose evaluation will produce the original value. Thus `o.lift()` produces an AST object `e` whose subsequent compilation and evaluation will produce `o`.

4 DSLs in XJ

This section describes two DSLs and shows how they can be implemented in XJ.

4.1 Processing Transaction Records

This example is due to Martin Fowler [16]. Suppose we have data that contains information about customer events. A customer event might be a service call including the name of the customer, the type of the call (failed hardware, billing query etc) and the date. This information may be provided in real-time or in a log file as text:

```
SVCLFLOWLER    10101MS0120050313
SVCLH0HPE     10201DX0320050315
SVCLTWO       x10301MRP220050329
USGE10301TWO  x50214..7050329
```

A Java program is to process the information. Obviously the first task of the Java application is to split the input strings up in terms of their fields. If there are a very small number of types of call then it would be OK to just write the appropriate string manipulation calls on the input. However, it would be better to define a string processing framework and use that. Fowler gives a framework-based implementation something like that shown in figure 4.

A DSL could be defined that makes the code much easier to read:

```
@Reader CallReader
map(SVCL,ServiceCall)
    4-18:CustomerName
    19-23:CustomerID
    24-27:CallTypeCode
    28-35:DataOfCallString
end
map(USGE,Usage)
    4-8:CustomerID
    9-22:CustomerName
    30-30:Cycle
    31-36:ReadDate
end
do
    ServiceCall
```

```

public void ConfigureCallReader(Framework framework) {
    framework.registerStrategy(ConfigureServiceCall());
    framework.registerStrategy(ConfigureUsage());
}

private ReaderStrategy ConfigureServiceCall() {
    ReaderStrategy result = new ReaderStrategy("SVCL",typeof(ServiceCall));
    result.addFieldExtractor(4,18,"CustomerName");
    result.addFieldExtractor(19,23,"CustomerID");
    result.addFieldExtractor(24,27,"CalltypeCode");
    result.addFieldExtractor(28,35,"DataOfCallString");
}

private ReaderStrategy ConfigureUsage() {
    ReaderStrategy result = new ReaderStrategy("USGE",typeof(Usage));
    result.addFieldExtractor(4,8,"CustomerID");
    result.addFieldExtractor(9,22,"CustomerName");
    result.addFieldExtractor(30,30,"Cycle");
    result.addFieldExtractor(31,36,"ReadDate");
}

```

Fig. 4: Configuring a Framework for Service Calls

```

Usage
end

```

which is equivalent to a class definition:

```

public class Callreader {
    public void ConfigureCallReader(Framework f) {
        // configuration calls...
    }
    private ReaderStrategy ConfigureServiceCall() {
        // Construction of service call strategy...
    }
    // More configurations...
}

```

The syntax-class for the Reader construct is defined in section 6.

4.2 Enterprise Java

Annotations are a new features of Java 5 which allow additional data to be embedded in program code in the form of name value bindings. Since annotations are a means of extending the Java language for specific use cases, they can be seen as an approach to defining Domain Specific variants of Java. An example of this is Enterprise Java where annotations are used to add in enterprise deployment information which is not facilitated in the standard Java language. Figure 5 is an example of defining a Calculator session bean which has a local and remote interface.

Although the annotations add in the necessary information, they arguably make the existing code more

difficult to read. If a new language syntax was designed for representing Enterprise Java information it would be significantly more concise and have a simpler mapping between the users' conceptual view of the program and its definition. Figure 6 shows the same program using an XJ defined syntax (given in Appendix 6.2):

A further example of applying XJ to Enterprise Java is in easing the specification of the object to relational mapping in entity beans. Figure 7 shows an example of a mapping between an order entity bean and its relational storage. Using the XJ technology to define a domain language, this could be expressed significantly more concisely as demonstrated in figure 8. Within this the persistence mapping is captured as part of the object's definition and the availability of getters and setters is expressed as a modifiers on the attribute definitions.

```

@Stateless
@Local ({Calculator.class})
@LocalBinding (jndiBinding="LocalCalculator")
@Remote ({RemoteCalculator.class})
@RemoteBinding (jndiBinding="RemoteCalculator")
public class LocalRemoteCalculator implements Calculator, RemoteCalculator {

    public double calculate (int start, int end, double growthrate, double saving) {
        double tmp = Math.pow(1. + growthrate / 12., 12. * (end - start) + 1);
        return saving * 12. * (tmp - 1) / growthrate;
    }

    public String getServerInfo () {
        return "This is the JBoss EJB 3.0 TrailBlazer";
    }
}

public interface Calculator {

    public double calculate (int start, int end, double growthrate, double saving);
}

public interface RemoteCalculator {

    public double calculate (int start, int end, double growthrate, double saving);

    public String getServerInfo ();
}

```

Fig. 5: Calculator Session Bean

```

import language Enterprise.EnterpriseBean;
import language Enterprise.RemoteLocal;
import language Enterprise.Remote;

@EnterpriseBean Calculator isStateless {

    @RemoteLocal double calculate(int start, int end, double growthrate, double saving) {
        double tmp = Math.pow(1. + growthrate / 12., 12. * (end - start) + 1);
        return saving * 12. * (tmp - 1) / growthrate;
    }

    @Remote String getServerInfo() {
        return "This is the JBoss EJB 3.0 TrailBlazer";
    }
}

```

Fig. 6: Calculator Session Bean in XJ

```
@Entity
@Table(name="ORDER_TABLE")
public class Order {

    private int id;
    private String address;

    @Id
    @Column(name="ORDER_ID")
    public int getId() {
        return id;
    }

    public void setId(int id) {
        this.id = id;
    }

    @Column(name="SHIPPING_ADDRESS")
    public String getAddress() {
        return address;
    }

    public void setAddress(String address) {
        this.address = address;
    }
}
```

Fig. 7: Order Entity Bean

```
import language Enterprise.EntityBean;

@EntityBean Order persistAs "ORDER_TABLE" {
    private int id persistAs "ORDER_ID" (get,set);
    private int address persistAs "SHIPPING_ADDRESS" (get,set);
}
```

Fig. 8: Order Entity Bean in XJ

5 Analysis

This paper has reviewed the motivation for language oriented programming (LOP) and domain specific languages (DSLs). It has described and analysed a number of approaches and technologies for LOP and DSLs and has made some recommendations regarding how LOP and DSLs can become part of mainstream system development.

The key recommendation is that the technology for LOP should become a standard part of a mainstream language and we have shown how this can be achieved by defining eXtensible Java (XJ). XJ supports LOP through the introduction of syntax-classes.

A syntax-class is a standard Java class extended with a grammar. Once defined, a syntax-class can be used as a construct within Java programs through the use of the '@' language escape character. LOP is also supported within XJ through standard access to Java abstract syntax and the use of quasi-quotes.

There are other systems that support LOP within Java. OpenJava (OJ) [9] supports LOP using a meta object protocol (MOP) that supports syntax expansion. Each class may specify a meta-class that is responsible for expanding its definition. The meta-class uses a standard definition interface to process its instances. This is a very flexible approach which supports complete access to abstract syntax. However, OJ does not include a mechanism for extending the concrete syntax of the language.

Maya [10] uses a pattern matching mechanism to define macro expansion rules on Java abstract syntax trees. The macros are defined using a new language concept called a *mayan* which defines a grammar production rule in terms of patterns. Maya allows macros to be overloaded based on the type of the arguments, thereby allowing the macros to expand differently. Maya offers similar features to XJ, however the changes to the language are global and not modularized by being attached to classes. The use of the syntax escape @-character, means that new syntax constructs are encapsulated, attached to classes and do not have global effect on the Java grammar. This is in contrast to Mayans that are independent of classes and have global scope much like macros in Lisp and Scheme.

The Java Syntactic Expander (JSE) [11] is a macro system that supports the definition of macros in Java and uses a quasi-quote mechanism to process the abstract syntax. JSE does not support new syntax rules and operate on the concrete syntax of the language. As we have seen in XJ it is sometimes necessary for syntax-classes to have raw access to AST structures.

Meta-AspectJ (MAJ) [12] is an annotation-based language extension mechanism that generates AspectJ. MAJ does not support concrete language extension.

The Jakarta Tool Suite (JTS) [13] provides a macro

facility and uses a quasi-quote like mechanism for dealing with abstract syntax. The expansion mechanism provides access to environments that are similar to the compiler contexts of XJ. JTS allows the concrete syntax of Java to be extended via the Bali parser. JTS is aimed at writing GenVoca generators and as such, although it provides a number of key LOP features, it is a pre-processor and therefore not as integrated with Java as the proposal for XJ.

The Fortress language [17] is an example of a new language that has included features that support LOP. The support is fairly simple: delimiters can be specified and the compiler supplies the raw string between the delimiters to user code for processing. However, this is an indication that LOP and DSLs are starting to emerge in languages aimed at the mainstream.

XJ is a proposal for a language extension to Java. In order for XJ to achieve its potential as a LOP system to support mainstream DSL development it must become part of the Java standard. A requirement for any proposal for mainstream language extension is that it should undergo a rigorous validation process. Although XJ has not been implemented in Java, it is one of the key features of the XMF language [15, 14] that has been used in commercial tools (XMF-Mosaic) and has been made open-source in 2008. All of the language mechanisms and algorithms necessary to implement XJ have been validated through the XMF system.

References

- [1] Language Oriented Programming. Martin Ward. Software - Concepts and Tools, Vol.15, No.4, pp 147-161, 1994
- [2] Steele, G. L. 1998. Growing a language. In Addendum To the 1998 Proceedings of the Conference on Object-Oriented Programming, Systems, Languages, and Applications (Addendum) (Vancouver, British Columbia, Canada). J. Haungs, Ed. OOPSLA '98 Addendum. ACM Press, New York
- [3] Martin Fowler's blog on Domain Specific Languages <http://www.martinfowler.com/articles/languageWorkbench.html>
- [4] XMF. Available from <http://www.ceteva.com>.
- [5] JEQUEL A Java embedded Query language (SQL in Java). <http://jequel.de/index.php?n=Main.JEQUEL>
- [6] Evolving an Embedded Domain Specific Language in Java OOPSLA 2006.
- [7] Architecture as Language: A Story. Marcus Voelter. <http://www.infoq.com/articles/architecture-as-language-a-story>.

- [8] Drools. <http://www.jboss.org/drools/>
- [9] OpenJava: A Class-based Macro System for Java. Michiaki Tatsubori, Shigeru Chiba, Marc-Olivier Killijian and Kozo Itano Lecture Notes in Computer Science 1826, Reflection and Software Engineering, , Walter Cazzola, Robert J. Stroud, Francesco Tisato (Eds.), Springer-Verlag, pp.117-133, 2000
- [10] Maya: Multiple-Dispatch Syntax Extension in Java. Jason Baker, Wilson C. Hsieh. The 2002 Conference on Programming Language Design and Implementation.
- [11] The Java Syntactic Extender (JSE). Jonathan Bachrach, Keith Playford. OOPSLA 2001.
- [12] Easy Language Extension with Meta-AspectJ. Shan Shan Huang, Yannis Smaragdakis. ICSE 2006.
- [13] JTS: tools for implementing domain-specific languages. Batory, D. Lofaso, B. Smaragdakis, Y. Proceedings. Fifth International Conference on Software Reuse, 1998. Publication Date: 2-5 Jun 1998 On page(s): 143-153.
- [14] Applied MetaModelling: A Foundation for Language Driven Development. Second Edition, 2008. Tony Clark, Paul Sammut, James Willans. E-book available from <http://www.ceteva.com>.
- [15] Superlanguages: Developing Languages and Applications with XMF. First Edition 2008. Tony Clark, Paul Sammut, James Willans. E-book available from <http://www.ceteva.com>.
- [16] DSLs. Martin Fowler presentation at JA00 2006. <http://www.infoq.com/presentations/domain-specific-languages>.
- [17] The Fortress Language Specification. Allen et al. Version 1.0 2008. Available from <http://research.sun.com/projects/plrg/Publications/fortress.1.0.pdf>

6 Appendix

```
import language java.syntax.Grammar;

public class Reader extends Sugar {

    private String name;
    private Vector<Mapping> mappings;
    private Strategy strategy;

    public Reader(String n,Vector<Mapping> m,Strategy s) {
        name = n;
        mappings = m;
        strategy = s;
    }

    @Grammar {
        Reader ::= n = Name '{' M = Mapping* s = Strategy '}'
            { new Reader(n,s,M) }.
        Mapping ::= 'map' '(' t = Name ',' n = Name ')' '{' F = Field* '}'
            { new Mapping(t,n,F) }.
        Field ::= s = Int '-' e = Int ':' n = Name
            { new Field(s,e,n) }.
        Strategy ::= 'do' N = Name*
            { new Strategy(@Cmp new MappingRef(n) | n <- N end) }.
    }
}
```

Fig. 9: Reader Syntax Class (Part 1 of 2)

```

public AST desugar(Context context) {
  return
  [| public class <name> {
    public void <"Configure" + name>(Framework framework) {
      <@Iterate Strategy s in strategies with e = [| |] {
        String name = "Configure" + s.getName();
        return
        [| <e>; framework.registerStrategy(<name>()); |]
      }>
      <@$Cmp([| private ReaderStrategy <"Configure" + m.getName() {
        ReaderStrategy result = new ReaderStrategy(t,typeof(<v>));
        <@Iterate Field f in m.getFields() with e = [| |] {
          AST start = f.getStartPos().lift();
          AST end = f.getEndPos().lift();
          AST name = f.getName().lift();
          return
          [| <e>; result.addFieldExtractor(<start>,<end>,<name>) |] ;
        }>
      } |]) {
        AST t = m.getTag().lift();
        AST v = Var(m.getName());
        Mapping m <- mappings;
      }$>
    }
  }
  |];
}
}
}

```

Fig. 10: Reader Syntax Class (Part 2 of 2)

```
import language java.syntax.Grammar;

public class EnterpriseBean Calculator extends Sugar {

    private String name;
    private boolean stateless;
    private AST signature;
    private Block body;

    public EnterpriseBean(String name,boolean stateless,AST signature,Block body) {
        this.name = name;
        this.stateless = stateless;
        this.signature = signature;
        this.block = block;
    }

    @Grammar extends Command {
        EnterpriseBean ::=
            name = Name
            stateless = 'isStateless' {true} | {false}
            signature = ClassSignature
            body = ClassBody {
                return new EnterpriseBean(name,stateless,signature,body);
            }
    }
}
```

Fig. 11: Syntax Class for Enterprise Beans (Part 1 of 3)

```

public AST desugar(Context context) {
    CompilationUnits cus = new CompilationUnits();
    cus.setContext(context);
    cus.add(createLocalInterface(context));
    cus.add(createRemoteInterface(context));
    cus.add(createBeanClass);
    return cus;
}

private CompilationUnit createBeanClass(Context context) {
    CompilationUnit bean = new CompilationUnit();
    bean.setContext(context);
    ClassDef beanDef = new ClassDef(name);
    signature.addImplements(name + "Local");
    signature.addImplements(name + "Remote");
    beanDef.setSignature(signature);
    if(stateless)
        bean.addAnnotation("Stateless");
    bean.add(beanDef);
    return bean;
}

private CompilationUnit createLocalInterface(Context context) {
    CompilationUnit local = new CompilationUnit();
    local.setContext(context);
    InterfaceDef localInterface = new InterfaceDef(name + "Local");
    Vector<MethodDef> methods = body.getMethodDefs();
    for(int i=0;i<methods.size();i++) {
        if(methoddef.hasProperty("Local") ||
            methoddef.hasProperty("RemoteLocal"));
        remoteInterface.add(local.getDeclaration());
    }
    local.add(localInterface);
    return local;
}

private void createRemoteInterface(Context context) {
    CompilationUnit remote = new CompilationUnit();
    remote.setContext(context);
    InterfaceDef remoteInterface = new InterfaceDef(name + "Remote");
    Vector<MethodDef> methods = body.getMethodDefs();
    for(int i=0;i<methods.size();i++) {
        if(methoddef.hasProperty("Remote") ||
            methoddef.hasProperty("RemoteLocal"))
            remoteInterface.add(remote.getDeclaration())
    }
    remote.add(remoteInterface);
    cus.add(remote);
}
}
}

```

Fig. 12: Syntax Class for Enterprise Beans (Part 2 of 3)

```
import language java.syntax.Grammar;

public class RemoteLocal extends Sugar {

    @Grammar extends Command {
        RemoteLocal ::= methodDef = MethodDef {
            methodDef.setProperty("RemoteLocal","true");
            return methodDef;
        }
    }
}

import language java.syntax.Grammar;

public class Remote {

    @Grammar extends Command {
        RemoteLocal ::= methodDef = MethodDef {
            methodDef.setProperty("Remote","true");
            return methodDef;
        }
    }
}
```

Fig. 13: Syntax Class for Enterprise Beans (Part 3 of 3)