This is a e-reader friendly version of the paper

Created May 1, 2011

Granularity in
Software Product Lines

Christian Kästner
School of Computer Science
University of Magdeburg, Germany
ckaestne@ovgu.de

Sven Apel
Department of Informatics and Mathematics
University of Passau, Germany
apel@uni-passau.de

Martin Kuhlemann
School of Computer Science
University of Magdeburg, Germany
mkuhlema@ovgu.de

Abstract

Building *software product lines (SPLs)* with features is a challenging task. Many SPL implementations support features with coarse granularity – e.g., the ability to add and wrap entire methods. However, fine-grained extensions, like adding a statement in the middle of a method, either require intricate workarounds or obfuscate the base code with annotations. Though many SPLs can and have been implemented with the coarse granularity of existing approaches, fine-grained extensions are essential when extracting features from legacy applications. Furthermore, also some existing SPLs could benefit from fine-grained extensions to reduce code replication or improve readability. In this paper, we analyze the effects of feature granularity in SPLs and present a tool, called *Colored IDE (CIDE)*, that allows features to implement coarse-grained and fine-grained extensions in a concise way. In two case studies, we show how CIDE simplifies SPL development compared to traditional approaches.

Categories and Subject Descriptors: D.2.3 [**Software**]: Software Engineering—*Coding Tools and Techniques*; D.2.6 [**Software**]: Software Engineering—*Programming*

1 Introduction

A *software product line (SPL)* aims at generating tailored programs from a set of features. Each feature represents an increment in functionality relevant to stakeholders. Different programs tailored for a given task or environment can be created by selecting a particular subset of features. Using SPLs it is possible to create program families of related programs for a domain.

A feature's implementation extends the program in one or more places. In prior work, we successfully created SPLs, e.g., (Apel and Batory, 2006; Kästner et al., 2007a; Leich et al., 2005), with several different implementation approaches like AHEAD (Batory et al., 2004), mixin layers (Smaragdakis and Batory, 2002), aspectual feature modules (Apel et al., 2008), or aspects (Kiczales et al., 1997). At that time, we have not been aware of the importance of granularity on feature implementation.

In a recent case study, we decomposed a legacy application – *Berkeley DB*, an embedded database engine[1] – into features in order to transform it into an SPL (Kästner et al., 2007a). We factored out 38 features, so that users could configure Berkeley DB to optionally

---

[1] http://www.oracle.com/database/berkeley-db

include features such as *transactions*, *statistics*, or *logging*. We noticed that this decomposition was much more difficult than our prior SPL implementations. Several features were hard to implement because they needed to introduce statements in the middle of methods, or add parameters to existing methods. The modularization of features became very difficult to understand and maintain. Though some problems were caused by limitations of AspectJ, we found that other languages like Jak (Batory et al., 2004), Jiazzi (McDirmid et al., 2001), or Hyper/J (Tarr et al., 1999) also poorly support extensions of such fine granularity.

What we observed is that languages that are based on *collaboration-based design* (Reenskaug et al., 1992) focus on adding new members to existing classes, new classes, and extending existing methods (Apel et al., 2007). These capabilities are sufficient for many applications. But for decomposing Berkeley DB, we needed much more fine-grained extensions. For example, when adding trans-action and synchronization mechanisms to the database system, we extended the program in over 240 places, often in places not trivially to access, like in the middle of a method (Kästner et al., 2007a). To realize such extensions, we had to use workarounds that obfuscated the code. In the end, we even considered restarting from scratch us-

ing preprocessor directives such as *'#ifdef'* and *'#endif'* to surround feature-specific code, but refrained because of other problems this would have introduced (Spencer and Collyer, 1992), which we discuss later.

Fine-grained extensions pose a major challenge for current approaches to feature-based SPL development, especially when creating a SPL by decomposing a legacy application. Alerted by these results we analyzed the granularity of extensions in SPLs we created earlier. Even though fine-grained extensions were required less frequently than in Berkeley DB, we found potential for improvements where current implementations replicate code or use workarounds that obfuscate the source code. We did not notice this potential earlier because we accepted the limitations of the given languages.

In this paper, we explore effects of granularity of different approaches on SPL development. We show that existing approaches are not able to implement fine-grained extensions satisfactorily and analyze possible solutions that allow implementing SPLs without sacrificing understandability. We present a tool, called *Colored IDE (CIDE)*, that combines the strengths of existing approaches and adds support to overcome the granularity problem. Finally, we illustrate the benefits of CIDE in two case studies, the

initial Berkeley DB case study and a small, existing SPL, in which we did not realize granularity effects at first.

2 Background: Approaches to SPL Implementation

There are two common ways to implement an SPL: the *compositional approach* and the *annotative approach*.

2.1 The Compositional Approach

Compositional approaches implement features as distinct modules. To generate a product line member, a set of modules is *composed*, usually at compile-time or deploy-time. There is a large body of work on feature composition usually employing component technologies (Szyperski, 2002), or specialized architectures and languages like frameworks (Johnson and Foote, 1988), mixin layers (Smaragdakis and Batory, 2002), AHEAD (Batory et al., 2004), multi-dimensional separation of concerns (Tarr et al., 1999), and aspects (Kiczales et al., 1997).

In Figure 1, we show an example of 3 modules implemented with Jak (AHEAD Tool Suite) (Batory et al., 2004),

a compositional approach, that uses mixin technologies. The first module (Lines 1–5) implements the basic stack. The other two modules implement extensions for the two features locking and logging. In Jak a feature can extend an existing class using the *refines* keyword and introduce new methods or extend existing ones with wrappers (by overriding using the *Super* keyword to call the original implementation).

2.2 The Annotative Approach

Annotative approaches implement features with some form of explicit or implicit annotations of the source code. Typical examples of explicit annotations are *'#ifdef'* and *'#endif'* statements of C/C++ style preprocessors to surround feature code. Others use Java annotations or new language constructs in the code to be extended. Examples are *explicit programming* (Bryant et al., 2002), *Frames/XVCL* (Jarzabek et al., 2003), *Spoon* (Pawlak, 2006), *Gears* (Krueger, 2002), *software plans* (Coppit et al., 2007), *metaprogramming with traits* (Turon and Reppy, 2007), and *aspects using annotations* (Kiczales and Mezini, 2005b). Alternatively, implicit annotations exploit existing language facilities. For example, deliberately introduced empty methods can be

```
1  class Stack {
2    boolean push(Object o) {
3      elementData[size++] = o;
4    }
5  }
```

```
6  refines class Stack {
7    boolean push(Object o) {
8      Lock l=lock(o);
9      Super.push(o);
10     l.unlock();
11   }
12   Lock lock(Object o) { /*...*/ }
13 }
```

```
14 refines class Stack {
15   boolean push(Object o) {
16     Super.push(o);
17     log("added " + o);
18   }
19   void log(String msg) { /*...*/ }
20 }
```

Figure 1: A basic stack and two features implemented in Jak.

used as hooks for extensions, or naming conventions can be employed for extensions like "synchronize all methods starting with *'sync_'*". In some approaches, features are partly modularized leaving just annotations in the code, at which feature code is later introduced, while in others the whole SPL including alternative and mutually exclusive features are encoded in a single code base and configurations are created by removing or deactivating code fragments.

Especially common is the use of preprocessors to configure a program. For example, with the C implementation of Berkeley DB a user can configure 11 different features at compile-time. An example code fragment from Berkeley DB with *'#ifdef'* statements for the features *HAVE_QUEUE* and *DIAGNOSTIC* is shown in Figure 2.

Most SPL implementations can be categorized into compositional and annotative approaches. Additionally, some approaches like aspects can be included in both groups, depending on how they are used, or combine approaches from both groups. However, there are techniques that do not fit into either group, e.g., those based solely on tool support (Beuche et al., 2004), based on *model-driven development* (Trujillo et al., 2007), or on *generative programming* (Czarnecki and Eisenecker, 2000). Still, in this paper we focus on compositional and annotative

```
1  static int __rep_queue_filedone(dbenv, rep, rfp)
2    DB_ENV *dbenv;
3    REP *rep;
4    __rep_fileinfo_args *rfp; {
5  #ifndef HAVE_QUEUE
6    COMPQUIET(rep, NULL);
7    COMPQUIET(rfp, NULL);
8    return (__db_no_queue_am(dbenv));
9  #else
10   db_pgno_t first, last;
11   u_int32_t flags;
12   int empty, ret, t_ret;
13 #ifdef DIAGNOSTIC
14   DB_MSGBUF mb;
15 #endif
16   // over 100 lines of additional code
17 }
18 #endif
```

Figure 2: Code excerpt of Berkeley DB.

approaches because they cover a significant fraction of contemporary SPL implementations and reveal interesting differences concerning granularity of extensions.

3 Granularity

Features can extend a program with additional code. Extensions with coarse granularity add new classes or methods to the program or extend explicit extension points. Plug-in architectures in frameworks and component approaches are typical approaches that provide possibilities for such extensions. However, this might not be sufficient. Developers might want to introduce new statements into existing methods and extend expressions or even method signatures. These are fine-grained extensions and require additional support. For example, they arise when legacy applications are feature-refactored (Liu et al., 2006). (These applications were not designed with feature granularity in mind, and consequently fine-grained extensions are commonly needed when extracting features). We evaluate possible levels of granularity and typical problems for both compositional and annotative approaches.

```
1  class Stack {
2    void push(Object o, Transaction txn) {
3      if (o==null || txn==null) return;
4      Lock l=txn.lock(o);
5      elementData[size++] = o;
6      l.unlock();
7      fireStackChanged();
8    }
9  }
```

Figure 3: Fine-grained extension example.

3.1 Compositional Approaches

Existing compositional techniques usually allow coarse-grained extensions only. Many provide mechanisms to define explicit extension points, which are the only points that can be extended by a feature. Others like AHEAD and AspectJ can extend virtually every method of the system. In most approaches it is only possible to introduce new classes, methods, or fields and to extend whole methods using wrappers (a.k.a. around advice, method refinements, or method overriding) (Apel et al., 2007).

We (Kästner et al., 2007a) and others (Murphy et al., 2001) noticed several limitations when implementing fine-grained extensions compositionally. We exemplify the three

most common limitations by means of the code snippet in Figure 3, in which the underlined code belongs to a synchronization feature and should be implemented in its own module.

1. **Statement Extensions.** In most compositional approaches it is not possible to introduce statements in the middle of an existing method in order to extend certain statements or sequences of statements therein.[2] For example, consider how to synchronize only the statement in Line 5. Simple wrappers around the whole method are not sufficient. Instead, we have to introduce the locking statements in Lines 4 and 6 specifically. Note, statement extensions might also access local variables. Usually, workarounds borrow from annotative approaches and introduce artificial extensions points for extension. Typically, a developer would introduce calls to empty hook methods (Murphy et al., 2001) or perform an Extract Method refactoring (Fowler, 1999) that moves Line 5 to its own method so that it can

---

[2]A notable exception is *AspectJ* that enables to extend method calls or field access inside specific methods (Kiczales et al., 2001; Kästner et al., 2007a). This feature can be used to emulate statement extensions in some cases (cf. Sec. 5.1).

be extended with a wrapper. Local variables, if accessed in the extension, are passed as parameters. Either workaround requires explicit or implicit annotations and severely obfuscates the source code (Murphy et al., 2001; Kästner et al., 2007a).

2. **Expression Extensions.** Extensions to an individual expression can occur as well. An example is shown in Line 3, in which the condition of the *if* statement is extended. A typical workaround again creates a new method and moves the expression there, so that it can be extended with wrappers.

3. **Signature Changes.** To the best of our knowledge, there is no compositional approach that allows to introduce an additional parameter into an existing method signature, as the *txn* parameter in Line 2. Instead, method signatures are considered unchangeable. Typical workarounds store the additional parameters in thread-safe fields, duplicate code, or use complex language mechanisms like the Wormhole Pattern in AspectJ (Laddad, 2003). However, all of these workarounds introduce different problems and reduce code quality (Rosenmüller et al., 2007; Kästner, 2007). Note that it is also necessary to adapt all calls to the extended method.

A possible implementation of the base code and its extensions using Jak is shown in Figure 4. Statements and expression extensions are implemented with the two hook methods *h1* and *h2*. The parameter is passed with a thread-safe field *pushTxn* and the original *push* method is deactivated by throwing an exception. Apparently, the extension code in Figure 4 is much larger than the amount of underlined code in Figure 3.

Extending compositional approaches with new language constructs for fine-grained extensions is not trivial either because of several conceptual problems. Firstly, signatures are used to identify the methods that are to be extended. If changing method signatures for an optional feature was possible, another naming scheme would need to be used to identify methods. Consequently, most languages consider signatures as immutable, and do not account for the possibility of signature changes.

Secondly, compositional approaches only introduce new code fragments in positions in which the order does not matter. Thus, it is possible to introduce new classes into the program or new methods into a class, but not new statements at a fixed position inside a method. This target position is not known when implementing the feature and could move if other features introduced statements as well.

```
1  class Stack {
2    void push(Object o) {
3      if (o==null || h1()) return;
4      h2(o);
5      fireStackChanged();
6    }
7    boolean h1() { return false; }
8    void h2(Object o) {
9      elementData[size++] = o;
10   }
11 }
```

```
12 refines class Stack {
13   ThreadLocal<Transaction> pushTxn =
           new ThreadLocal<Transaction>();
14   void push(Object o, Transaction txn) {
15     pushTxn.set(txn);
16     Super.push(o);
17   }
18   void push(Object o) {
19     throw new UnsupportedOperationException(
20         "Call push(Object,Transaction) instead");
21   }
22   boolean h1() {
23     return pushTxn.get() == null;
24   }
25   void h2(Object o) {
26     Lock l = pushTxn.get().lock(o);
27     Super.h2(o);
28     l.unlock();
29   }
30 }
```

Figure 4: Fine-grained extension with AHEAD.

18

Therefore, compositional approaches usually offer only wrappers that add statements at the beginning and/or the end of a method, but not at a finer granularity. Similarly, parameters in method signatures are ordered, which makes parameter introductions difficult.

The coarse granularity of compositional approaches leads to several problems when developers use them to implement fine-grained extensions anyway. Workarounds often obfuscate the source code and are verbose and hard to understand. Many workarounds replicate code or use heavy-weight architectures that induce performance penalties (Murphy et al., 2001; Kästner et al., 2007a; Rosenmüller et al., 2007; Kästner, 2007).

3.2 Annotative Approaches

Conceptually, annotations can mark code fragments at arbitrary levels of granularity. They simply introduce markers at the exact positions that should be extended. Typical examples are C/C++ style preprocessors, which although they annotate only whole physical lines, are sufficient for even the finest extensions due to the ability to isolate language constructs in separate lines. A possible preprocessor implementation of our example from Figure 3 is shown in

```
1  class Stack {
2    void push(Object o
3  #ifdef TXN
4    , Transaction txn
5  #endif
6    ) {
7      if (o==null
8  #ifdef TXN
9        || txn==null
10 #endif
11       ) return;
12 #ifdef TXN
13     Lock l=txn.lock(o);
14 #endif
15     elementData[size++] = o;
16 #ifdef TXN
17     l.unlock();
18 #endif
19     fireStackChanged();
20   }}
```

Figure 5: Fine-grained extension with C/C++ preprocessor.

Figure 5. Other annotation approaches allow similarly fine-grained extensions, e.g., (Jarzabek and Shubiao, 2003; Coppit et al., 2007).

Annotations do not share the conceptual limitations regarding ordered statements and fixed signatures because they indicate the final position in the base code. Therefore,

20

a method can always be identified by its final signature and also the position of a statement or a parameter in an ordered list can always be determined.

Still, our experience[3] and reports by others, e.g., (Spencer and Collyer, 1992; Baxter and Mehlich, 2001), show that annotations have problems as well. Firstly, annotations themselves obfuscate the source code as apparent in Figure 5. Secondly, annotating arbitrary code fragments, whether they make sense or not, is problematic. For example, it is possible to annotate an opening bracket with one feature and the closing bracket with another. This makes annotations error-prone and raises complexity. Thirdly, there may be problems dealing with separating terminals like commas between parameters. There are frequent situations when such simple syntactic elements must be annotated for features as well.

In Figure 6 we show an example: an *init* method with two parameters, in which the first parameter is included only if *transactions* are enabled in the system, and the second is included only if *logging* is enabled. However, when annotating this code fragment with C/C++ style preproces-

```
1  void init(Transaction txn,
       LoggingLevel level){/*impl.*/}
```

```
2  void init(
3  #ifdef TRANSACTION
4    Transaction txn
5  #ifdef LOGGING
6    ,
7  #endif
8  #endif
9  #ifdef LOGGING
10   LoggingLevel level
11 #endif
12   ){/*impl.*/}
```

Figure 6: Decomposition with preprocessors.

sors we have to split the method declaration into multiple
lines and even include the comma inside a nested *'#ifdef'*
statement so that all derivable variants are syntactically
correct.

Despite these problems, annotative approaches sup-
port fine-grained extensions better than compositional ap-
proaches. However, they provide no perceptible form of
modularity.

4 CIDE

4.1 Overview

Motivated by the problems of both compositional and annotative approaches, we built an Eclipse-based prototype tool for decomposing legacy applications into features that may have a fine granularity.[4] It uses the semantics of preprocessors, i.e., it can be classified as annotative approach, but avoids the pollution of source code. Developers start with a fully composed application with all features implemented in a single code base, typically a legacy application. Then, they make features explicit by successively associating code fragments with one or more features, i.e., they mark the corresponding code. Alternatively, developers can also extend the application with new features and associate all new code fragments with those features. Just as with *'#ifdef'* statements, code fragments are only included when all associated features are selected in a given configuration.

In contrast to traditional preprocessors as in C/C++, we do not obfuscate the source code with additional an-

---

[4]The tool can be downloaded at http://wwwiti.cs.uni-magdeburg.de/iti_db/research/cide.

notations. Instead, we use the representation layer of the editor to indicate the associated features with different background colors. Thus, developers can directly recognize whether a code fragment is associated with a feature. In case a code fragment is associated with multiple features, which is traditionally done with nested preprocessor statements, we mix the according background colors (e.g., red + blue = purple). Feature names are shown in tool-tips on request. Note that it is usually not possible to recognize the features of a code fragment solely by background colors, especially when many features overlap. However, colors are sufficient to determine the beginning and the end of a code fragment associated with a set of features, and it is convenient to look up the actual features using tool-tips. Because of its colorful appearance (cf. Fig. 7) we named the tool *Colored Integrated Development Environment (CIDE)*.

As with preprocessors it is still possible to insert or edit code, while the colors remain assigned to the code fragments. But even though CIDE is based on preprocessor semantics, we do not assign features to arbitrary code fragments to avoid the problems of meaningless associations and syntactical elements illustrated in Section 3.2. Instead, we assign features to *structural code elements*
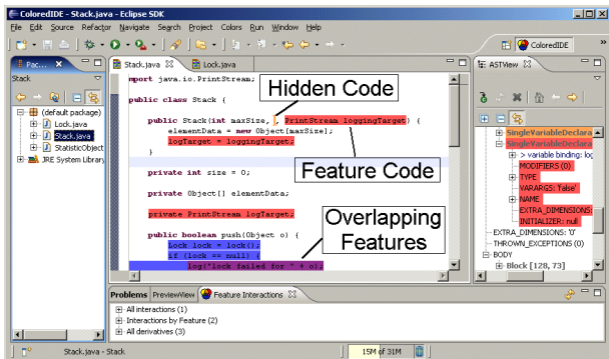
Figure 7: CIDE Screenshot.

of the source code. The form of such structural elements depends on the artifact type. For example, programs in many languages can be expressed as *abstract syntax tree (AST)*, not only Java programs. In all these cases we use this underlying structure and associate features to structural elements.

4.2 Coping with Feature Granularity

We illustrate the mechanisms used in CIDE by the example of Java code and its AST representation. An AST contains all structural elements that are relevant for the source code. It is possible to recreate the source code from the AST (except its original formatting). In Figure 8, we exemplify a small Java code snippet and a corresponding AST. Note, the AST does not include syntactical elements like commas or brackets.

Instead of using offset and length to specify code fragments, we assign features directly to subtrees of the AST as defined by the language grammar. For example, we can assign a feature to a class node and all its children (i.e., making a class optional), or to a statement and its children as highlighted in Figure 8 (i.e., statements are optional child nodes of *block* nodes). It is even possible
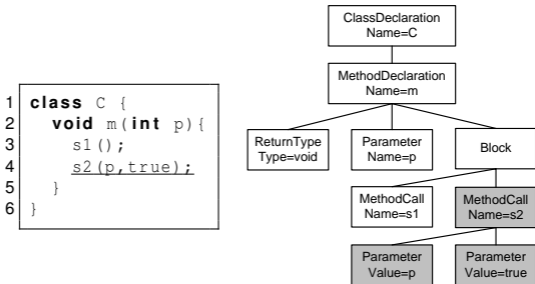
Figure 8: AST Example.

to assign a single parameter node to a feature (i.e., an optional child node of a method node). However, nodes that are not optional in the AST (i.e., that cannot be removed without invalidating the AST), like the return type of a method cannot be associated with features individually. Developers can but are not required to view the code as AST, usually, in the editor they assign the selected code fragments directly to features, which CIDE automatically maps to AST elements.

For Java code, we carefully defined two kinds of exceptions from the subtree rule exemplified in Figure 9. Firstly, statements that embrace other statements can be associated with features without necessarily associating its child elements. So, it is possible to remove individual *if*, *for*, *while*, *try*, and similar statements without removing the statements they surround. Secondly, children of binary expression nodes can be associated with features even though they are not optional AST elements (every binary expression node requires two child expressions). This allows us do decompose expression statements as in Figure 3. We decided to include these exceptions for increased granularity.

The restriction of assigning only optional AST subtrees to features enforces decomposition into 'reasonable' code

```
1  void m(boolean
        a, boolean
        b) {
2    a = a || s1()
         || b;
3    if (a) {
4      s2();
5    }
6  }
```
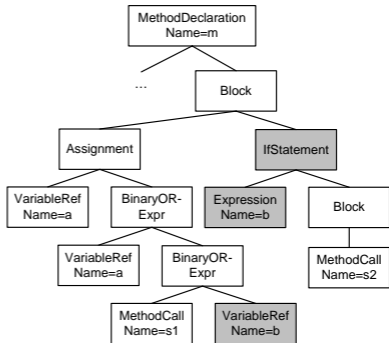
Figure 9: Exceptions to the optional subtree rule.

fragments. Thus, developers can no longer associate only an opening bracket but not the closing one with a feature. Furthermore, they do not need to deal with syntactical elements such as commas in Figure 6 at all. We allow only AST operations that create a well-formed AST as output, so the Java compiler can *parse* the generated code of any configuration. At the same time, the AST structure is extremely flexible, allowing developers to associate arbitrary statements or even parameters or parts of expressions with features. Decomposing a code snippet like the one introduced in Figure 3 is straightforward, because parameters, statements and expressions are all AST nodes that can simply be 'colored'.

Note, the granularity of CIDE is not as fine as that provided by C/C++ like preprocessors. For example, developers cannot specify two alternative result types of a method, because the AST only allows one child element for the result type per method. In practice we have not yet found any cases where this was a limitation. Furthermore, for reasons of simplicity developers can assign only a set of features to a code fragment and not a propositional formula (Czarnecki and Pietroszek, 2006). Consequently, it is not possible to specify a code fragment to be included if a feature is *not* selected (*#ifndef*). However, we are able

30

to work around these problems by modeling features differently in a feature model (e.g., specify two alternative features instead of a single optional one). If future case studies showed the demand for such expressibility, we could still extend CIDE at the cost of a more complicated user interface.

4.3 Feature Management

Feature implementations with fine-grained extensions are difficult to understand in both compositional and annotative techniques, as illustrated in Section 3. Even though compositional approaches can implement features cohesively, workarounds make it hard to understand extensions. Also the feature code in annotative approaches is hard to understand because it is scattered and annotations obfuscate the source code.

CIDE enhances feature management (navigation, selection, composition) with tool support. Firstly, navigation support allows developers to jump between code fragments associated with a certain feature set to diminish the problems of missing feature cohesion. This makes it easy to find all code fragments associated with a feature. Because feature code is still placed where it extends the

program, and it is therefore obvious to see how it extends the program, it is simple to understand how a feature is implemented (cf. Fig. 3). The navigation support is in line with prior approaches of *concern graphs* (Robillard and Murphy, 2002) and the *map metaphor* (Griswold et al., 2001) that choose tool support over complex modular implementations.

Secondly, a projection facility can hide all code associated with a feature in the editor during development, so that the remaining code can be viewed in isolation. The projection facility is implemented like code folding in modern development environments, in which bodies of methods or comments can be folded and unfolded on request. When the developer requests a feature to be hidden, CIDE just leaves a marker to indicate hidden code. Thus, the developer can focus on selected features and hide feature code that is not relevant to the current task. The markers are still useful for modifications as they alert the developer of feature extensions that might be necessary to adapt as well.

Finally, we implemented a mechanism to export the marked code into cohesive modules implemented with compositional approaches (Kästner et al., 2007b). Our prototype currently provides exports to Jak and to AspectJ.

This way, feature implementations can be used independently of CIDE, e.g., for further code generation and optimization steps.

We argue that, compared to existing compositional and annotative approaches, CIDE makes it easier to develop feature-based SPLs because it hides all programming language concepts of refinements, aspects, hook methods, annotations, and similar for implementing features that are necessary in these approaches. It provides different views on the source code. Developers can hide features, navigate between them, and even export them. Even though CIDE cannot replace cohesive modules, it supports understanding scattered feature implementations and enables developers to choose the best view on the source code for each task.

5 Case Studies

In order to support our proposal, we conducted two case studies. Firstly, we analyze our decomposition of the embedded database engine Berkeley DB, a large scale legacy application with 84,000 LOC. Secondly, we study a small SPL of graph algorithms with 2,000 LOC designed as an

SPL using a compositional approach. In both, we compare their implementations with CIDE. The code for both case studies is available for download from the CIDE web page.[5]

5.1 Berkeley DB

Our initial insight that the granularity of existing compositional approaches was insufficient came from a case study, in which we decomposed a large scale legacy application - the embedded database engine Berkeley DB JE – into 38 features to make features explicit in its design and to make it configurable as an SPL. To decompose Berkeley DB, we first located and marked feature code, then removed this feature code, and finally reintroduced it using feature modules implemented with AspectJ (Kiczales et al., 2001), a compositional approach. Due to space restrictions, we limit our report to statistics and the most important insights in anecdotical form.

Berkeley DB was not designed as an SPL with features in mind. Instead it was implemented as a single application with a clean object-oriented design modularized in classes and packages. Some configurability was achieved using external parameters and runtime checks. Consequently,

---

feature code was scattered all over the application. Some larger features like Synchronization and Transaction affected up to 30 (of 300) classes in Berkeley DB in over 150 places (cf. (Kästner, 2007; Kästner et al., 2007a) for detailed statistics).

Decomposition with AspectJ. When decomposing Berkeley DB we were confronted with fine-grained extensions in almost every feature. For example, feature code was often located right in the middle of a method. Of 1,144 extensions used to implement the 38 features, 640 extensions (56 %) introduced new classes, methods, or fields. 214 other extensions (19 %) were simple method extensions that added wrappers to existing methods. They were well supported by AspectJ. However, 261 extensions (23 %) were required at statement level and 24 (2 %) at expression level, which posed major problems.

Statement extensions were implemented in two different ways. Firstly, AspectJ supports extensions finer than just wrappers around methods: it is possible to extend methods *calls* or field accesses inside specific methods by combining *call*, *set*, or *get* with *withincode* pointcut designators. We used such extensions to emulate statement extensions when the feature code was placed directly before, after, or around a single method call or field access.

This was sufficient for 121 extensions (46 % of all statement extensions). However, our solution was rather fragile, because any changes to the internal implementation of the base code might require the pointcuts to be altered. For the remaining 140 statement and 24 expression extensions we had to introduce hook methods and extend those.

We also faced the problem that certain method parameters of the code base belonged to a feature and should be removed when the feature is not selected. This made it very difficult to detach some features in Berkeley DB, because AspectJ, like other compositional approaches, does not permit changes to method signatures. We first noticed the problem when we decomposed the part of the transaction system that is responsible for atomicity. To group several operations as a single transaction, the user can create a *Transaction* object and pass it to operations like *put*, *get*, or *delete*. This object is then passed further along the control flow to acquire suitable locks. Inside Berkeley DB, such a parameter is included in the signature of 59 methods. Without the transaction feature, none of these methods should include a transaction parameter.

We experimented with different possibilities to implement the Transaction feature. We removed the parameter from the base code and introduced a second method for

each operation with the additional parameter that calls the first method. We could pass the parameter in a thread-protected field, use method objects (Fowler, 1999), use dummy or default parameters, employ AspectJ's wormhole pattern (Laddad, 2003), or use parameter objects that encapsulate all parameters. Unfortunately, all options obfuscate the source code. Firstly, excited by AspectJ's advanced constructs, we implemented the wormhole pattern in 16 pieces of advice as exemplified in Figure 10. There, all calls to *getWritebleLocker(Environment)* are replaced by calls to *getWriteableLocker(Environment, Transaction)*, and the transaction parameter is captured from interface methods like *openDatabase* and *get* using the *cflow* pointcut designator. Eventually, we found this too fragile and too hard to understand (Kästner et al., 2007a). Finally, we settled with parameter objects named *OperationContext* and interface methods that created those objects for internal use. To extend the method, we only needed to extend the parameter objects.

While we could work around the parameter limitation for the transaction feature, we did not attempt any further decomposition of features, in which a large amount of parameters were involved. Especially the Locking feature appeared unmanageable with compositional approaches

```
1  pointcut interfaceCallWithTransaction(Transaction
       txn) :
2      (execution(* Environment.openDatabase(..)) &&
           args(txn,*,*)) ||
3      (execution(* Database.get(..)) &&
           args(txn,*,*,*)) || ...;
4  pointcut getWritableLocker(Environment env):
       call(Locker LockerFactory.getWritableLocker(Environme
       && args(env);
5
6  Locker around(Environment env, Database db,
       Transaction txn) throws DatabaseException :
7      getWritableLocker(env) && this(db) &&
           cflow(interfaceCallWithTransaction(txn))
8      { return LockerFactory.getWritableLocker(env,
           txn); }
```

Figure 10: Using the Wormhole Pattern to pass a parameter.

as it had to introduce parameters like *locker*, *lockMode*, or *lockType* 289 times. Decomposing the Locking feature would either result in utterly unreadable code or require a complete preliminary redesign of the whole database engine.

Together, workarounds for fine-grained extensions of statements, expressions, and parameter lists made the decomposition very difficult and the implementation tedious.

Though it was possible to decompose most features, code quality suffered and the resulting code base became hard to understand and maintain.

Decomposition with CIDE. After these discouraging initial results, we experimented with other compositional approaches and with decomposing Berkeley DB using a C/C++ preprocessor but limitations discussed in Section 3.2 held us from trying them on a case study of this size. Finally, we decomposed Berkeley DB once more into the same features using CIDE. In CIDE, fine-grained extensions were no problem, because we could directly associate single statements, exceptions, or even parameters with features. This avoids all previously required workarounds. Because we did not need to perform any manual decomposition or implementation of feature modules, but only assign features to code fragments inside the IDE, the decomposition of Berkeley DB was more convenient and much faster (3 days instead of 1 month)[6]. CIDE's finer granularity supported decomposing Berkeley DB significantly. We were even able to decompose the Locking feature with its 289 parameter extensions.

---

[6]The second decomposition was also faster because we were already familiar with the source code. Still, the difference is significant.

By using CIDE, understanding the SPL became simpler, because neither do additional statements obfuscate the source code, nor is it necessary to understand complex workarounds like in Figures 4 and 10. Moreover, the projection facilities helped to reason about features in isolation or in concert with other features.

5.2 The Graph Product Line

The *graph product line (GPL)* is a small SPL of graph algorithms with about 2,000 LOC. It was designed from scratch as an SPL and suggested as a benchmark for SPL technologies (Lopez-Herrejon and Batory, 2001). The domain of graphs was chosen because it is well understood and the algorithms are well known in computer science. It consists of 14 features: edges can be directed, undirected and optionally weighted, there are two search methods breadth first search (BFS) and depth first search (DFS), and finally there are several graph algorithms like cycle checking, shortest path, strongly connected, or minimum spanning tree. The implementation as an SPL makes it possible to select the graph properties and algorithms needed for a given problem (Lopez-Herrejon and Batory, 2001). The original version of GPL was implemented with

Figure 11: UML class model of GPL (excerpt).

mixin layers (Smaragdakis and Batory, 2002), but other implementations e.g. in AspectJ (Lopez-Herrejon and Batory, 2002a) and Hyper/J (Lopez-Herrejon and Batory, 2002b) are available. All these implementations use compositional approaches with similar results regarding the following analysis.

Although the GPL case study and the figures below appear almost trivial, we intentionally chose it because it is well suited to show that there are benefits from implementation approaches that support fine-grained extensions, *even* for SPLs that are (1) not extracted from legacy applications but designed from scratch and (2) small and simple.

In Figure 11 we show the basic class structure of GPL. Details of some features are omitted. The architecture is

simple: a *Graph* class holds the graph consisting of edges and vertices. Most features add a method to the *Graph* class and extend the *run* method used for benchmarks. Some features like *Number*, *Shortest*, or *Weighted* add fields for temporary values to *Vertex* and *Edge* and extend their *display* method to show the additional values. Features like *Connected* or *Cycle* also depend on methods implemented in previous features like search algorithms.

Implementation with Mixin Layers. The granularity required to implement GPL is coarse. As illustrated in Table 1 most features introduce only new code fragments (25 methods and 21 fields into existing classes and 18 new classes). Methods were rarely extended (by means of method overriding). Of 84 methods in GPL only 4 were ever extended: the method *Vertex.display* was extended by 9 features, the benchmark method *Graph.run* was extended by 7 features and *Edge.display* and *Edge.adjust* were extended once. Not a single hook method was used to implement extensions at statement or expression level.

Even though all features could be implemented with the coarse granularity provided by the language, we found some code replication that could have been avoided if fine-grained extensions were available. Firstly, the mutually exclusive search features *BFS* and *DFS*, as well as

```
1   public void search(Workspace w) {
2     int s = vertices.size();
3     if (s == 0) return;
4     for (int c = 0; c < s; c++)
5       vertices.get(c).init_vertex(w);
6     for (int c = 0; c < s; c++) {
7       Vertex v = vertices.get(c);
8       if (!v.visited) {
9         w.nextRegionAction(v);
10        v.dfSearch(w);
11        v.bfSearch(w);
12      }
13    }
14  }
```

Figure 12: Search method differs only in a single line.

Directed and *Undirected*, introduced similar methods that differ only in minor parts. Alternative implementations could use hook methods inside a common base feature to avoid code replication. In CIDE both variants can be combined in a single code base, in which both features only extend a single line as shown in Figure 12 (feature code is italicized and underlined respectively).

Secondly, the extension of the *Weighted* feature introduces a second *addAnEdge* method and a second *Edge* constructor with an additional parameter, though concep-

tually the *addAnEdge* method and the constructor of the *Edge* class should be both extended with an additional *weight* parameter. So, with the *Weighted* feature enabled there were two *addAnEdge* methods, one with two and one with three parameters, even though only the method with three parameters should be used by clients.

Implementation with CIDE. In CIDE, we have composed the whole code base including several mutually exclusive features in a single code base, in which different code fragments are associated with features. This way, the file containing the code of the class *Graph* became fairly large (772 LOC), of which 88 % were associated with features. The projection facility came in handy to see only relevant code fragments.

We were surprised that despite GPL's apparent coarse granularity, we could still identify several situations, in which we benefit from CIDE's fine granularity and could reduce code replication. For example, we integrated both *search* methods that only differ in one line as shown in Figure 12. Similar code replication between the *Directed* and *Undirected* features could be avoided as well. Furthermore, we could implement the additional parameter for the *Weighted* feature much simpler without additional methods or constructors as shown in Figure 13. We sim-

```
1  class Edge {
2    Edge(Vertex start, Vertex end, int weight) {
3      this.start = start; this.end = end;
4      this.weight = weight;
5    }
6  }
7  class Graph {
8    Edge addAnEdge(Vertex start, Vertex end,
         int weight) {
9      Edge theEdge = new Edge(start, end, weight);
10     edges.add(theEdge);
11     start.addNeighbor(new Neighbor(end, theEdge));
12     return theEdge;
13   }
14 }
```

Figure 13: Parameter extension for addAnEdge.

ply associated all *weight* parameters with the *Weighted* feature (underlined). Because of reduced code replication and because we do not need implementation overhead of mixin layers, the code size of GPL in CIDE is 36 % smaller than in the original implementation (1,222 LOC instead 1,920 LOC).

6 Related Work

SPL Adoption. Krueger distinguishes between three different adoption models for SPLs (Krueger, 2002). Firstly, the *proactive adoption model* relates to the waterfall approach of conventional software engineering. The SPL is planned up front and designed with all features in mind, therefore few carefully planned extension points are often sufficient and fine-grained extensions arise infrequently. Due to their coarse granularity, compositional approaches are well suited, though in our GPL case study we still discovered possibilities for improvement.

In contrast, the *extractive adoption model* converts an existing legacy applications to an SPL by decomposing its features – aiming to reduce adoption time, cost and risk. In legacy applications, features were not planned in the design phase and consequently dissolved in the implementation. Redesigning the complete legacy application is usually out of the question, because the additional effort could exceed the gained benefit of reduced time, costs, and risk. Extracting scattered feature implementations requires the ability to implement fine-grained extensions, so that annotative approaches and especially CIDE are better suited for SPL implementation. The Berkeley DB case

study illustrates the need for fine-grained extensions when extracting features.

Finally, the *reactive adoption model*, which relates to spiral software engineering in conventional software, stands between the proactive and extractive adoption models. Not all features are foreseen and designed from scratch, but the SPL is extended in iterative steps. In this model, most extensions are preplanned, but it is not possible to design for extensibility in every case. So, we expect a medium to high number of fine-granular extensions as well.

This emphasizes the relevance of fine-grained extensions and the applicability of CIDE in SPLs beyond only decomposing legacy applications.

Related Tool Approaches. The concept of using the representation layer to show additional information without obfuscating the source code as in CIDE was already used by several development environments. Recent examples are *presentation extension* (Eisenberg and Kiczales, 2007), *AspectBrowser* (Griswold et al., 2001) and *Spotlight* (Coppit et al., 2007).

IDEs for visual programming and intentional programming abstract from traditional code representations and store code in internal tree formats close to ASTs. The

idea of storing program code in databases to allow flexible queries to create different views on the code goes back to Linton (Linton, 1984). Modern examples are *Snippets* (Westphal et al., 2004), *effective views* (Janzen and De Volder, 2004), and the *Domain Workbench* (Simonyi et al., 2006), that store all code in internal tree structures, similar to how features are assigned to code in CIDE.

There is also an impressive body of work on feature annotations of source code. However, such work usually does not aim at SPLs but at *virtual separation of concerns* to make concerns or features explicit. For example, Robillard and Murphy suggested *concern graphs* where developers can collect methods belonging to a feature in an external window (Robillard and Murphy, 2002). Work on *visual separation of concerns (VSC)* extends this and provides aggregated views on the source code by features (Chu-Carroll et al., 2003). Furthermore, the *AspectBrowser* (Griswold et al., 2001) and *JQuery* (Janzen and De Volder, 2003) use pattern expressions or queries to find code fragments belonging to a certain feature. Similarly to CIDE, feature annotations are stored externally, the code itself is not changed. While *FEAT*, *ConcernMapper*, *VSC*, and *JQuery* work at method level, thus providing only coarse granularity, *AspectBrowser* works on character level of unparsed

code. In contrast to CIDE, these tools are used exclusively for source code exploration and navigation, whereas CIDE is designed as a software product line tool suite, that uses additionally annotations to configure the program.

Closest to our work on SPLs is an extended UML modeling environment suggested by Czarnecki and Pietroszek, in which users can assign individual model elements to features (Czarnecki and Pietroszek, 2006). A tool uses this information to check whether all possible configurations are well-formed. However, while CIDE assigns only one or more features to elements, they can assign arbitrary propositional formulas like 'entity X is only included if feature X is selected or feature Y is not'. Compared to CIDE, this raises the complexity for both the user interface and the back-end.

Cross-Section Views. Projections on features in CIDE resemble cross-section views, e.g., in 3D engineering or tomography. They hide all details unnecessary for the current task and let the user focus on certain details from different views. An early example of such cross-section views in software engineering is the concept of *program slicing* (Weiser, 1984), in which the source code is projected to the fragments relevant for a certain control flow. Typically, it shows all code that can affect a selected variable

and hides all other code. Program slicing helps abstracting from the whole program and to focus on a concrete, usually comprehension or maintenance focused task. Instead of projecting by control flow, CIDE projects (structurally) on individual features. This enables developers to understand programs by features and supports switching between different views dynamically.

Also the notion of on-demand remodularization follows the vision of cross-section views on a program. On-demand remodularization as introduced by Ossher and Tarr is the ability to extract a concern of an application into a new module without affection other concerns (Ossher and Tarr, 2000). Typically, it is desired to decompose a program in one dimension first, e.g., data and classes, and later extract a concern in another dimension, e.g., features, which was not considered in the first decomposition. On-demand remodularization permits developers to identify and encapsulate new concerns at any time without necessity to rearchitect the software. This enables the developer to create only those modularizations that they need and as they need them (Ossher and Tarr, 2000).

Additionally to the virtual separation of concerns approaches introduced above, there are several approaches based on composition languages and tool support that

physically remodularize the source code. An example is *effective views* (Janzen and De Volder, 2004) that are however limited to two dimensions, classes and 'modules'. Though modules are similar to features, they may not overlap. Another proposed tool for supporting remodularization in multiple dimensions is the *concern manipulation environment* that is aimed at describing and extract concerns (Harrison et al., 2005). Similarly, CIDE is a tool-based approach and allows to incrementally decompose a system into further features by associating code fragments with (potentially overlapping) features. Features can also be extracted in later development steps when needed. The key difference however is that CIDE aims directly at SPLs.

Feature Cohesion and Modular Reasoning. Implementing features modularly is a fundamental goal of compositional approaches. The motivation is that developers are able to understand and modify features in isolation without resorting to global reasoning or affecting other parts of the system if they are implemented cohesively (Stevens et al., 1974). There is a large body of work that discusses modular reasoning for different programming paradigms and languages, e.g., (Stata and Guttag, 1995; Kiczales and Mezini, 2005a). Furthermore, Tarr et al. postulate that it should be possible to decompose a system in different di-

mensions and still implement all concerns modularly (Tarr et al., 1999). However, in our observation, such cohesive feature modules tend to become so complex that the benefit of modular reasoning diminishes in the presence of fine-grained extensions and with current compositional approaches.

In contrast, annotative approaches like preprocessors do not implement features cohesively. On the contrary, feature implementations are scattered throughout the system, preventing modular reasoning. To understand a feature, developers must search the feature code as a first step. However, approaches for virtual separation of concerns listed above like *concern graphs* support identifying and reasoning about scattered code. Although CIDE still provides exports into feature modules, it follows these ideas and provides tools to reason about scattered feature implementations instead of enforcing complex cohesive implementations.

Feature Model and Consistency. CIDE, as most other approaches for implementing extensions for features, does not deal directly with consistency constraints or use a feature model. A feature model on top of CIDE that describes the relationships between features, e.g., that one feature depends on another feature, is necessary for a holistic

approach, but not relevant for the granularity discussion in this paper. Once features are related to each other, CIDE can check the feature association of all AST elements to ensure that every configuration is not only parseable, but also compilable. This requires that several constraints are fulfilled, e.g., that the target methods of all method calls are defined in all configurations. These checks follow the concepts of Czarnecki and Pietroszek (Czarnecki and Pietroszek, 2006) and Thaker et al. (Thaker et al., 2007) but their discussion exceeds the scope of this paper.

7 Conclusion

There are many ways to implement features. However, when features have fine-grained extensions, as common when decomposing a legacy application, their implementations tend to become complicated, unreadable, and unmaintainable. Compositional approaches do not support fine-grained extensions, so that workarounds are required which raise the implementation's complexity. In contrast, annotative approaches can implement fine-grained extensions but introduce readability problems by obfuscating the source code.

To avoid these problems when developing SPLs with fine-grained extensions in a concise fashion, we built a tool, called CIDE, that simplifies SPL development. It is based on preprocessor semantics but uses background colors instead of source code statements and by providing the possibility to hide features it avoids obfuscating the code. CIDE restricts features to structural code elements in order to simplify usage for developers while still providing fine granularity. Finally, CIDE supports developers in understanding features with navigation and projection facilities and the possibility to export the SPL into distinct feature modules.

In two case studies we showed the advantage of CIDE over existing compositional approaches. It was possible to implement fine-grained features including statement and expression extensions and even signature changes without resorting to workarounds. With CIDE we could implement additional features that were not reasonably possible with compositional approaches and reduce code replication of earlier implementations.

In ongoing and future work, we add support for additional artifact types (e.g., C, C#, JavaScript, Grammars, XML) and formalize the criteria when an element can be colored. We intend to extend our export facility with ad-

ditional target languages to analyze modular reasoning. Further, we plan to extend CIDE as a round-trip engineering tool where it is possible to edit the exported code and reimport it back to CIDE. This will eventually enable us to use CIDE to provide additional views on existing SPLs.

References

Apel, S., and Batory, D. (2006). When to use features and aspects? a case study. In *GPCE*.

Apel, S., Leich, T., and Saake, G. (2008). Aspectual feature modules. In *IEEE TSE*. Online first.

Apel, S., Lengauer, C., Batory, D., Möller, B., and Kästner, C. (2007). An Algebra for Feature-Oriented Software Development. Tech. Rep. MIP-0706, University of Passau.

Batory, D., Sarvela, J. N., and Rauschmayer, A. (2004).

Scaling step-wise refinement. *IEEE TSE*, *30*(6), 355–371.

Baxter, I., and Mehlich, M. (2001). Preprocessor conditional removal by simple partial evaluation. In *Proc. Working Conference on Reverse Engineering*. Washington, DC, USA: IEEE Computer Society.

Beuche, D., Papajewski, H., and Schröder-Preikschat, W. (2004). Variability management with feature models. *Sci. Comput. Program.*, *53*(3), 333–352.

Bryant, A., et al. (2002). Explicit programming. In *AOSD*, pp. 10–18. New York, NY, USA: ACM Press.

Chu-Carroll, M., Wright, J., and Ying, A. (2003). Visual separation of concerns through multidimensional program storage. In *AOSD*, pp. 188–197. New York, NY, USA: ACM Press.

Coppit, D., Painter, R., and Revelle, M. (2007). Spotlight: A prototype tool for software plans. In *ICSE*, pp. 754–757. Washington, DC, USA: IEEE Computer Society.

Czarnecki, K., and Eisenecker, U. (2000). *Generative programming: methods, tools, and applications*. ACM Press.

Czarnecki, K., and Pietroszek, K. (2006). Verifying feature-based model templates against well-formedness ocl constraints. In *GPCE*, pp. 211–220. New York, NY, USA: ACM Press.

Eisenberg, A., and Kiczales, G. (2007). Expressive programs through presentation extension. In *AOSD*, pp. 73–84. New York, NY, USA: ACM Press.

Fowler, M. (1999). *Refactoring. Improving the Design of Existing Code*. Addison-Wesley.

Griswold, W. G., Yuan, J. J., and Kato, Y. (2001). Exploiting the map metaphor in a tool for software evolution. In *ICSE*, pp. 265–274. IEEE Computer Society.

Harrison, W., et al. (2005). Concern modeling in the concern manipulation environment. In *ICSE Workshop on Modeling and Analysis of Concerns in Software*, pp. 1–5. New York, NY, USA: ACM Press.

Janzen, D., and De Volder, K. (2003). Navigating and querying code without getting lost. In *AOSD*, pp. 178–187. New York, NY, USA: ACM Press.

Janzen, D., and De Volder, K. (2004). Programming with

crosscutting effective views. In *ECOOP*, vol. 3086 of *Lecture Notes in Computer Science*, pp. 195–218. Springer.

Jarzabek, S., Bassett, P., Zhang, H., and Zhang, W. (2003). Xvcl: Xml-based variant configuration language. In *ICSE*, p. 810. Los Alamitos, CA, USA: IEEE Computer Society.

Jarzabek, S., and Shubiao, L. (2003). Eliminating redundancies with a "composition with adaptation" meta-programming technique. In *ESEC/FSE*, pp. 237–246. New York, NY, USA: ACM Press.

Johnson, R. E., and Foote, B. (1988). Designing reusable classes. *Journal of Object-Oriented Programming*, *1*(2).

Kästner, C. (2007). *Aspect-Oriented Refactoring of Berkeley DB*. Master's thesis, University of Magdeburg, Germany.

Kästner, C., Apel, S., and Batory, D. (2007a). A case study implementing features using aspectj. In *SPLC*.

Kästner, C., Kuhlemann, M., and Batory, D. (2007b). Automating feature-oriented refactoring of legacy applications. In *ECOOP Workshop on Refactoring Tools*.

Kiczales, G., and Mezini, M. (2005a). Aspect-oriented programming and modular reasoning. In *ICSE*, pp. 49–58. New York, NY, USA: ACM Press.

Kiczales, G., and Mezini, M. (2005b). Separation of concerns with procedures, annotations, advice and pointcuts. In A. Black (Ed.) *ECOOP*, vol. 3586 of *Lecture Notes in Computer Science*, pp. 195–213. Springer.

Kiczales, G., et al. (1997). Aspect-oriented programming. In M. Aksit, and S. Matsuoka (Eds.) *ECOOP*, vol. 1241 of *Lecture Notes in Computer Science*, pp. 220–242. Springer.

Kiczales, G., et al. (2001). An overview of aspectj. In J. L. Knudsen (Ed.) *ECOOP*, vol. 2072 of *Lecture Notes in Computer Science*, pp. 327–353. Springer.

Krueger, C. (2002). Easing the transition to software mass customization. In *Intl' Workshop on Software Product-Family Eng.*, pp. 282–293. London, UK: Springer-Verlag.

Laddad, R. (2003). *AspectJ in Action: Practical Aspect-Oriented Programming*. Manning Publications.

Leich, T., Apel, S., and Saake, G. (2005). Using step-wise refinement to build a flexible lightweight storage manager. In *ADBIS*, pp. 324–337.

Linton, M. (1984). Implementing relational views of programs. *SIGPLAN Not.*, *19*(5), 132–140.

Liu, J., Batory, D., and Lengauer, C. (2006). Feature oriented refactoring of legacy applications. In *ICSE*, pp. 112–121.

Lopez-Herrejon, R., and Batory, D. (2001). A standard problem for evaluating product-line methodologies. In *GCSE*, pp. 10–24. London, UK: Springer-Verlag.

Lopez-Herrejon, R., and Batory, D. (2002a). Using aspectj to implement product-lines: A case study. Tech. rep., Dept. of Computer Sc., University of Texas at Austin.

Lopez-Herrejon, R., and Batory, D. (2002b). Using hyper/j to implement product-lines: A case study. Tech. rep., Dept. of Computer Sc., University of Texas at Austin.

McDirmid, S., Flatt, M., and Hsieh, W. (2001). Jiazzi: New-age components for old-fashioned java. In *OOPSLA*, pp. 211–222.

Murphy, G., et al. (2001). Separating features in source code: an exploratory study. In *ICSE*, pp. 275–284. Washington, DC, USA: IEEE Computer Society.

Ossher, H., and Tarr, P. (2000). On the need for on-demand remodularization. In *ECOOP Workshop on Aspects and Dimensions of Concerns*.

Pawlak, R. (2006). Spoon: Compile-time annotation processing for middleware. *IEEE Distrib. Sys. Onl.*, *7*(11), 1.

Reenskaug, T., et al. (1992). Oorass: Seamless support for the creation and maintenance of object-oriented systems. *J. OO Programming*, *5*(6), 27–41.

Robillard, M., and Murphy, G. (2002). Concern graphs: Finding and describing concerns using structural program dependencies. In *ICSE*, pp. 406–416. New York, NY, USA: ACM Press.

Rosenmüller, M., Kuhlemann, M., Siegmund, N., and Schirmeier, H. (2007). Avoiding variability of method signatures in software product lines: A case study. In *GPCE Workshop on Aspect-Oriented Product Line Engineering*.

Simonyi, C., Christerson, M., and Clifford, S. (2006). Intentional software. In *OOPSLA*, pp. 451–464. New York, NY, USA: ACM Press.

Smaragdakis, Y., and Batory, D. (2002). Mixin layers: an object-oriented implementation technique for refinements and collaboration-based designs. *ACM TOSEM*, *11*(2), 215–255.

Spencer, H., and Collyer, G. (1992). #ifdef considered harmful or portability experience with C news. In *USENIX Conf.*, pp. 185–198.

Stata, R., and Guttag, J. (1995). Modular reasoning in the presence of subclassing. In *OOPSLA*, pp. 200–214. New York, NY, USA: ACM Press.

Stevens, W., Myers, G., and Constantine, L. (1974). Structured Design. *IBM Systems Journal*, *13*(2).

Szyperski, C. (2002). *Component Software: Beyond Object-Oriented Programming*. Addison-Wesley.

Tarr, P., et al. (1999). N degrees of separation: Multidimensional separation of concerns. In *ICSE*, pp. 107–119. Los Alamitos, CA, USA: IEEE Computer Society Press.

Thaker, S., et al. (2007). Safe composition of product lines. In *GPCE*.

Trujillo, S., Batory, D., and Diaz, O. (2007). Feature oriented model driven development: A case study for portlets. In *ICSE*, pp. 44–53. Washington, DC, USA: IEEE Computer Society.

Turon, A., and Reppy, J. (2007). Metaprogramming with traits. In *ECOOP*.

Weiser, M. (1984). Program slicing. *IEEE TSE*, *10*(4), 352–357.

Westphal, B., Harris, F., and Dascalu, S. (2004). Snippets: Support for drag-and-drop programming in the redwood environment. *Journal of Universal Computer Science*, *10*(7).

Feature	CI	CR	MI	Extended Methods
Prog	1	0	0	-
Benchmark	0	1	7	-
Directed	4	0	0	-
Undirected	4	0	0	-
BFS	2	1	3	Vertex.display
DFS	2	1	3	Vertex.display
Weighted	0	2	2	Edge.display, Edge.adjust
Transpose	0	1	1	-
Connected	1	2	1	Vertex.display, Graph.run
Cycle	1	2	1	Vertex.display, Graph.run
MSTKruskal	0	2	1	Vertex.display, Graph.run
MSTPrim	0	2	1	Vertex.display, Graph.run
Number	1	2	1	Vertex.display, Graph.run
Shortest	0	2	1	Vertex.display, Graph.run
StronglyC.	2	2	1	Vertex.display, Graph.run

CI: class introductions; CR: class refinements; MI: method introductions

Table 1: Feature Statistics for GPL.