

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

A novel model-based testing approach for software product lines

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1614384> since 2017-10-01T15:09:52Z

Published version:

DOI:10.1007/s10270-016-0516-2

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

This is an author version of the contribution:

Ferruccio Damiani, David Faitelson, Christoph Gladisch, Shmuel Tyszberowicz.

A novel model-based testing approach for software product lines.

Software & Systems Modeling.

DOI: 10.1007/s10270-016-0516-2

published on-line on Software & Systems Modeling (First Online: 13 February 2016)

The on-line published version is available at:

<http://link.springer.com/article/10.1007%2Fs10270-016-0516-2>

A novel model-based testing approach for software product lines

Ferruccio Damiani · David Faitelson ·
Christoph Gladisch · Shmuel Tyszberowicz

Received: date / Accepted: date

Abstract Model-based testing relies on a model of the system under test. FineFit is a framework for model-based testing of Java programs. In the FineFit approach, the model is expressed by a set of tables based on Parnas tables. A software product line is a family of programs (the products) with well-defined commonalities and variabilities that are developed by (re)using common artifacts. In this paper we address the issue of using the FineFit approach to support the development of correct software product lines. We specify a software product line as a specification product line where each product is a FineFit specification of the corresponding software product. The main challenge is to concisely specify the software product line while retaining the readability of the specification of a single system. To address this we used delta-oriented programming, a recently proposed flexible approach for implementing software product lines, and developed: (i) delta-tables as a means to apply the delta-oriented programming idea to the specification of software product lines; and (ii) DeltaFineFit as a novel model-based testing approach for software product lines.

The authors of this paper are listed in alphabetical order. This work has been partially supported by project HyVar (www.hyvar-project.eu), which has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 644298; by ICT COST Action IC1402 ARVI (www.cost-arvi.eu); by ICT COST Action IC1201 BETTY (www.behavioural-types.eu); by Italian MIUR PRIN 2010LHT4KM project CINA (sysma.imtlucca.it/cina); by Ateneo/CSP D16D15000360005 project RunVar; and by GIF (grant No. 1131-9.6/2011).

F. Damiani
University of Torino, Dipartimento di informatica, C.so Svizzera 185, 10149 Torino, Italy
(E-mail: ferruccio.damiani@unito.it)

D. Faitelson
Afeka Tel Aviv Academic College of Engineering, Israel (E-mail: davidf@afeka.ac.il)

C. Gladisch
Karlsruhe Institute of Technology, Germany (E-mail: gladisch@ira.uka.de)

S. Tyszberowicz
The Academic College of Tel Aviv Yaffo, Israel (E-mail: tyshbe@tau.ac.il)

Keywords Java · Alloy · Software product line · Delta-oriented programming · Model-based testing · Refinement

1 Introduction

In this paper we propose a novel model-based testing approach for *software product lines* (SPLs). An SPL is a set of programs (the products) that share significant common functionality and have well-understood and organized variabilities [9,33]. Our idea is to integrate data refinement based testing into SPL development.

FineFit [17] is an approach for model-based testing of Java programs which relies on the notion of *data refinement* [34] to compare the state of the model with the state of the *system under test* (SUT). Data refinement captures the relationship between an abstract model and its concrete representation. During the testing process, FineFit must both retrieve abstractions of the SUT current state and instruct the SUT to perform operations with specific input. These tasks are supported by two Java code fragments—the *retrieve function* and the *driver*, respectively. The retrieve function and the driver are written by FineFit’s users. A prototypical implementation of the FineFit tool is available [45].

In the FineFit approach the structure of the model and the specification of the system’s operations are given by a set of tables, based on Parnas tables [32]. Parnas tables organize expressions, where rows and columns separate an expression into cases, and each table entry specifies either the result value for some case or a condition that partially identifies some case. The strength of the tabular notation is its clear readability which helps in reducing errors in specifications.

In this paper we address the problem of using the FineFit approach to support the development of correct SPLs. Writing for each product to be tested a FineFit specification, a retrieve function, and a driver would be error prone and not cost-effective. Our idea is to express the *specification of an SPL* as a product line, where each product is a FineFit model (i.e., a set of tables) for the corresponding product (i.e., a Java program) of the specified SPL.

The main challenge addressed in this paper is to devise a means to concisely specify the SPL being tested while retaining the readability of a FineFit specification of a single system. To this aim, we consider *delta-oriented programming* (DOP) [36,7] (see also [1], Section 6.6.1), a recently proposed flexible approach for implementing SPLs. DOP is an extension of *feature-oriented programming* (FOP) [5], a prominent approach for developing SPLs (cf. [1], Section 6.1).¹ DeltaJ [7,22,44] is the archetypal language for delta-oriented programming of SPLs of Java programs.

As pointed out in [1], “much of the tremendous power of features is yet to be unlocked by making features explicit throughout the entire systems and software lifecycle”. In this paper we present DeltaFineFit, a novel model-based testing approach for delta-oriented SPLs. DeltaFineFit integrates data refinement based testing into delta-oriented SPL development by ensuring that each product is generated together with its FineFit model and the suitable driver and retrieve functions.

¹ A straightforward embedding of FOP into DOP is illustrated, e.g., in [38].

DeltaFineFit enables the fully automated testing of all the products of an SPL. When the number of products is too large, testing all the products is unfeasible. This could be addressed by using, e.g., *sample-based SPL testing* techniques [21, 20, 27, 23], where a subset of products—covering relevant combinations of features—is generated and tested by applying single system testing techniques.

The main contribution of this paper is to introduce the *delta table* concept: an approach for deriving a new table from an existing one, based on the difference (the delta) between the tables. Delta tables are used to define a notion of “delta module for FineFit specifications” (i.e., a construct that describes how to modify the FineFit specification of a product to obtain the FineFit specification of another product), that we call *delta-table module*. Delta tables are illustrated and evaluated (in terms of their support to conciseness and readability of SPL specifications) by considering a small SPL as a case study that is used as a running example throughout the paper.

A prototypical implementation of DeltaTables, a tool that generates a FineFit specification by applying a delta-table module to a FineFit specification, is available at the DeltaFineFit home page [43]. All the tables presented in the paper that are the result of delta-table application have been automatically generated. The DeltaFineFit tool chain (which will provide fully automated support to the integrated use of DeltaTables, DeltaJ, and FineFit) is currently under development.

The remainder of the paper is organized as follows. Section 2 recalls the FineFit approach for specifying and testing single Java programs. Section 3 recalls DOP and illustrates, by using DeltaJ, the SPL example that is used through the paper. Section 4 introduces delta tables. Section 5 illustrates DeltaFineFit by means of the SPL introduced in Section 3. Section 6 evaluates the DeltaFineFit approach. Section 7 reviews some related work. We conclude in Section 8 that also describes some future directions. Appendix A elaborates on the structure and the semantics of the FineFit tables. Appendix B illustrates the algorithm that describes the semantics of delta tables.

A preliminary version of the material presented in this paper is briefly outlined in [10]. Here we present a new and a more flexible notion of delta-tables, describe implementation details, and provide detailed explanations and examples.

2 A recollection of data-refinement testing of Java programs with FineFit

The data refinement theory [34] captures the relationship between an abstract model and its concrete implementation. The state spaces of the two levels of abstraction are related by a *retrieve function* which maps the concrete representation into the abstract one. When we use data refinement for testing, the abstract model becomes a test oracle and a source of test cases for the concrete program.

The FineFit [17, 45] model-based (or model-driven) testing framework uses data refinement to directly compare the state of the model (the specification) with the state of the SUT. FineFit helps to understand the testing results and to trace the reason for any difference between the specification and the SUT (see [17]). To test a system, the user writes its specification as a collection of HTML² tables (where cells contain

² In practice the tables can be written using any tool that can export its output to HTML, for example MS Word. FineFit ignores anything that is not part of an HTML table.

Alloy expressions [19]). The specification of the SUT consists of tables for its basic types (called *atoms*), for the states of the SUT, for its invariants, and one table for each operation provided by the SUT. Each operation table defines an operation as a predicate on the system states immediately before and after the operation (similar to how operations are specified in languages like Z [40]). FineFit uses these predicates in two ways. First it uses the predicates to test the behavior of the SUT’s operations. For each operation X it applies the corresponding predicate to the SUT’s state, right after the SUT completes operation X . A false result indicates a discrepancy between the expected and the actual behaviors. Second it uses the predicates to generate test cases by “solving” each predicate (a solution is an assignment of values that satisfy the predicate to the state variables and the inputs), and uses the solutions as test cases. These two capabilities are implemented with the help of Alloy’s³ Kodkod [41] relational model finder. We present an example in the rest of this section; for more details see [17].

The SUT’s retrieve function is a Java method that translates the concrete state (of the SUT) into an instance of the abstract state. The task of implementing the retrieve function is delegated to the developer—who must know how the data structures implement the system’s specification. This makes the testing framework flexible and scalable, since the programmers can control the parts (i.e., select a subset) of the system to expose for the purpose of testing, even if the system is too large or complicated for automatic analysis.

Figure 1 illustrates the testing process of a SUT when FineFit is used. Before the actual testing process is executed, FineFit checks the consistency of the model. Then FineFit begins the testing process by arbitrarily choosing an operation and input available in the current state, applying the corresponding concrete operation to the SUT, and checking whether the new state corresponds to the operation’s specification. This process continues until either a discrepancy is found or until the user decides to stop the process. During testing, FineFit prints a trace that consists of the abstract snapshots (states) and the operation calls of the SUT. When a problem is detected, the user can review the entire history that led to it.

In the following subsections we demonstrate data-refinement testing with FineFit by an example of a simple Java program implementing a photo album, that we call the Base Album. We begin by describing in Section 2.1 the program we would like to test. Then we describe in Section 2.2 the FineFit model that is used to test the program. Finally, in Section 2.3, we present the interfaces between FineFit and the program. These interfaces are used to retrieve the current program state for analysis in FineFit and to allow FineFit to apply operations on the program.

2.1 The Java program to be tested

Listing 1 shows the Java code implementing the Base Album. Informally, the Base Album manages a sequence of photos and provides operations for adding and view-

³ Alloy is a general purpose modeling language (in the style of Z) for reasoning about relational structures with first order logic. It has no direct concept of system states or operations, and it does not offer any tool for testing software.




Fig. 1 FineFit testing procedure. Each abstract state is a Kodkod Instance object that maps the names of the state variables to relations; each relation is a set of tuples of atoms. Each abstract activations consists of an operation name and atoms for the input parameters of the operation. The driver is explained in Section 2.3.

ing photos. It consists of the class `Photo`, the interface `PhotoAlbum` (with the nested classes `PhotoExists` and `AlbumIsFull`), and the class `ArrayPhotoAlbum` which implements the interface, providing functionality for adding a photo to the album and viewing its photos (the methods `addPhoto` and `viewPhotos` of `PhotoAlbum`, respectively).

2.2 The FineFit model

In order to test a system with FineFit, the developer has to write a model that specifies the desired behavior of the system. A FineFit model consists of relational variables and invariant predicates that represent the system states and of predicates that describe the effects of the operations on the system states. Both the structure (i.e., the state variables and the invariant predicates) of the model and its operations predicates are written in a tabular form. Formally a FineFit *model* (or *specification*) consists of:

1. a *constants table* that defines the values of constants used in the model;
2. an *atoms table*, listing the name and the *scope* (i.e., the maximum number of instances) of each of the *atoms* (i.e., the basic types) used in the model;
3. a number of *enumeration tables* listing the values of each enumerated type, where each value has a unique symbolic name;
4. a *state table*, listing the name and the type of the variables representing the state of the model;
5. an *invariants table*, listing the invariants that the state of the model must satisfy when the model is created and that must be preserved by any operation that may be executed on the model;
6. a number of *operation tables* representing the operations of the model.

```

1 public class Photo {
2     private String image; // represents the bitmap image
3     public Photo(String theImage) {
4         if(theImage == null) throw new IllegalArgumentException("NullImage");
5         image = theImage;
6     }
7     public String getImage() { return image; }
8     public String toString() { return image; }
9 }

1 import java.util.Set;
2 public interface PhotoAlbum {
3     Photo addPhoto(String image);
4     Set<Photo> viewPhotos();
5     public class PhotoExists extends RuntimeException { }
6     public class AlbumsFull extends RuntimeException { }
7 }

1 import java.util.Set;
2 import java.util.HashSet;
3 public class ArrayPhotoAlbum implements PhotoAlbum {
4     private int size = 0;
5     private Photo [] photoAt;
6     public ArrayPhotoAlbum(int maxSize) {
7         if(maxSize<1) throw new IllegalArgumentException("IllegalSize");
8         photoAt = new Photo[maxSize];
9     }
10    public boolean imageIsInAlbum(String image) {
11        for (int i= 0; i < size; i++) {
12            Photo p = photoAt[i];
13            if (p.getImage().equals(image)) return true;
14        }
15        return false;
16    }
17    public Photo addPhoto(String image) {
18        if(image == null) throw new IllegalArgumentException("NullImage");
19        if(size == photoAt.length) throw new AlbumsFull();
20        if(imageIsInAlbum(image)) throw new PhotoExists();
21        Photo new_photo = new Photo(image);
22        photoAt[size] = new_photo;
23        size = size + 1;
24        return new_photo;
25    }
26    public Set<Photo> viewPhotos() {
27        Set<Photo> result = new HashSet<Photo>();
28        for (int i= 0; i < size; i++) { result.add(photoAt[i]); }
29        return result;
30    }
31 }

```

Listing 1: Java 1.5 code implementing the Base Album

We name the first five kinds of tables *structural tables* (since they represent the structure of the model). The syntax of the tables of the model is illustrated in Fig. 2. The model of the Base Album is given in Fig. 3. The expressions used in the tables are written in Alloy [19], except that we use variables decorated with “?” and “!” to denote inputs and outputs, respectively. In most cases the expressions are self-explanatory; when needed we provide a short explanation.

Constants table syntax		Atoms table syntax		Enumeration table syntax	
Constant name	Value	Atom	Scope	Atom	
$Name_1$	$Value_1$	$Atom_1$	$Number_1$	$Value_1$	
\vdots	\vdots	\vdots	\vdots	\vdots	
$Name_{n_C}$	$Value_{n_C}$	$Atom_{n_A}$	$Number_{n_A}$	$Value_{n_E}$	

State table syntax		Invariants table syntax	
State variable	Type	Invariant name	Invariant
var_1	$Type_1$	$Name_1$	$AlloyFormula_1$
\vdots	\vdots	\vdots	\vdots
var_{n_V}	$Type_{n_V}$	$Name_{n_I}$	$AlloyFormula_{n_I}$

Operation table syntax						
$id(x?:Type, y!:Type)$	$C_1^{\{1..\}} \dots C_1^{\{..h..\}} \dots C_1^{\{..m\}}$					
	$\dots \dots \dots \dots \dots \dots$					
	$C_p^{\{1\}} \dots C_p^{\{h\}} \dots C_p^{\{m\}}$					
$varOrOut_1$	$t_1^1 \dots t_1^h \dots t_1^m$					
\vdots	$\vdots \dots \vdots \dots \vdots \dots \vdots$					
$varOrOut_n$	$t_n^1 \dots t_n^h \dots t_n^m$					

Fig. 2 Model tables syntax

2.2.1 Structural tables

We now describe the structural tables of the Base Album example.

Constants table. The constant table contains one constant, MAX, that specifies the maximal number of photos in the album.

Atoms (basic types) table. The atoms used in the model of Base Album are Photo, Int, and Report. Report is an enumeration type that we use in conjunction with an output variable to indicate the success or the failure of an operation (its values are specified by an enumeration table).

The atoms table associates a scope with each atom to determine its maximal number of instances used during testing. In the example we test the system with at most five photos. Note that the scope is not part of the specification; it is only used to limit the amount of entities in the test cases. The scope of the Report type can be safely set to three, because Report has exactly three instances—one for each constant in the enumeration table.

Enumeration table. The Base Album has an enumeration table for Report. It defines Report to be a set of exactly three instances whose names are given in the table.

Constants table		Atoms table		Enumeration table	
Constant name	Value	Atom	Scope	Report	
MAX	5	Photo	5	ALBUM_FULL	
		Int	5	PHOTO_EXISTS	
		Report	3	OK	

State table		Invariants table	
State variable	Type	Invariant name	Invariant
photoAt	seq Photo	Disjoint	#photoAt.elems = #photoAt
		BoundedSize	#photoAt ≤ MAX

Operation tables	
init()	true
photoAt	none -> none

addPhoto(p?:Photo, report!:Report)	#photoAt < MAX		#photoAt ≥ MAX
	p? !in ran[photoAt]	p? in ran[photoAt]	true
photoAt	photoAt.add[p?]	=	=
report!	OK	PHOTO_EXISTS	ALBUM_FULL

viewPhotos(result!: seq Photo)		true
result!		photoAt.elems

Fig. 3 Specification of the Base Album. Note that the expression `photoAt.add[p?]` is equivalent to the expression `photoAt+(#photoAt→ p?)`, that is, it describes the union of two relations. The content of the cells in the tables is always free of side-effects; there are no imperative statements in the FineFit language.

State table. The state table describes the state of the Base Album by capturing, via the variable `photoAt`, the sequence of photos in the album.

Invariants table. Two invariants are needed for the Base Album: (i) all the album photos are distinct (i.e., each photo has been added to the album only once), and (ii) there are no more than MAX photos in the album. The first invariant is specified using the number of photos in the album (recall that the album is a sequence of photos). This number has to be equal to the size of the set of all photos in the album.⁴

2.2.2 Operation tables

A FineFit operation table is a predicate that specifies the behavior of an operation as a relation between the model's state variables before the operation starts (the pre-state) and after the operation completes (the post-state). It consists of two major areas: a *precondition tree* which consists of *condition cells* (first p rows in the operation table syntax at the bottom of Fig. 2) and an *expression table* which consists of variable and value cells (last n rows in the operation table syntax). The precondition tree consists

⁴ Note that by definition a set does not contain duplicate elements.

of predicates that determine which columns to use in the definition of the post-state. The expression table is a set of columns, where each column defines the values of the state variables in the post state (given their values in the pre-state) and the values of the output parameters. For example, the `addPhoto` operation of the Base Album (cf. Fig. 3) updates the state of the `photoAt` state variable and the `report!` output parameter (first column). This operation has three expression columns—columns two to four in the table. The first expression column (the second column in the table) is for the case where the input photo `p?` (of type `Photo`) does not appear in the album and the album is not full, then the photo is added to `photoAt`. The next expression column is for the case where the input photo already appears in the non-full album (`photoAt` does not change), and the last expression column is for the case where the album is full (`photoAt` does not change). The value of `report!` (of type `Report`) is OK only in the case that the `addPhoto` operation indeed added a new photo to the album. Note that predicates that appear on top of each other are conjoined, while predicates that appear side by side represent disjunction. If we consider the precondition part of an operation table (see the operation table syntax, Fig. 2) as a matrix of p rows and m columns, then

- each of the predicates $C_p^{\{1\}} \dots C_p^{\{m\}}$ in row p occupies exactly one cell, and
- each of the predicates $C_q^{\{1..\}} \dots C_q^{\{..h..\}} \dots C_q^{\{..m\}}$ in row q (with $1 \leq q \leq p$) spans over one or more consecutive cells (a predicate that spans over columns $j..k$ is superscripted by the set $\{j..k\}$ of the indexes of those columns).

We provide further explanations about the structure and semantics of operation tables in Appendix A.

In addition to `addPhoto`, the model of the Base Album provides the operations `viewPhotos` and `init`. The operation `viewPhotos` uses `elems` to return (using `result!`) the set of elements in a given sequence (in our case the `photoAt` sequence). Finally, the operation `init`—which corresponds a class constructor in Java—defines the initial state of the album, demanding that there are no photos in the album.

The meaning of ‘=’ in a cell of the expression table part of an operation table (e.g., `addPhoto` in Fig. 3) is that the content of the state variable associated to the table row is left unchanged.

2.3 The retrieve function and the driver

During the testing process, the SUT must provide FineFit with abstractions of its current state, and FineFit must instruct the SUT to perform operations with specific inputs. The first task is supported by a *retrieve function* and the second by a *driver*.⁵

The retrieve function relates the concrete state of the SUT (Java code) to the abstract state of the specification. To provide a retrieve function the SUT must implement a method that returns an abstract snapshot (a collection of named sets of tuples, one for each state variable) of its current state. For example, in our Base Album application we have added the retrieve method illustrated in Listing 2 to the class

⁵ In the original paper about FineFit [17] we used the term *fixture*, which was borrowed from Fit [31], instead of *driver*. We think that driver is a more appropriate term.

```

1  import com.finefit.sut.IdMap;
2  import com.finefit.sut.State;

1  public State retrieve() {
2      State state = new State();
3      state.add_state("photoAt", 2);
4      for (int i = 0; i < size; i++) {
5          state.get_state("photoAt").add("" + i, IdMap.instance().obj2atom(photoAt[i]));
6      }
7      return state;
8  }

```

Listing 2: Required import statements and method implementing the retrieve function for the Base Album specification (to be inserted in the class `ArrayPhotoAlbum` in Listing 1)

`ArrayPhotoAlbum` (given in Listing 1). The abstract state is constructed using the `State` class that is part of the `FineFit` library. The class `IdMap` manages the mapping between the concrete Java objects and their corresponding abstract atoms.

A driver is needed to connect the SUT with `FineFit`. It is responsible for translating the abstract operation calls and data used by `FineFit` to the concrete representation of the SUT, and for translating concrete return values and exceptions back to their corresponding abstract atoms. To facilitate the implementation of the driver, we provide (as a part of the `FineFit` library) the base class `FineFitDriver`. It contains two maps, one (represented by the field `ops`) for translating abstract operation names to concrete operations and one (represented by the field `exceptions`) for translating concrete exceptions to abstract error codes. Any concrete driver class, like the `PhotoAlbumDriver` class for the Base Album application in Listing 3, has to populate these two maps with the exceptions and the operations of the SUT. This is done by defining the methods `set_up_operation_table` and `set_up_exception_table`, respectively. The concrete driver class has also to define the methods `set_up_sut` (for creating the SUT instance), `set_init_sut` (for performing further initialization operations that might be needed), and `retrieve` (for calling the `retrieve` method on the SUT). The driver unpacks the operation names and arguments coming from `FineFit`, calls the operation, captures the exceptions thrown by the operation, converts them into error codes and returns (by calling the SUT's `retrieve` method) to `FineFit` the state of the SUT after the operation has completed.

3 A recollection of delta-oriented programming of SPLs of Java programs with DeltaJ

`DeltaJ` is the archetypal language for delta-oriented programming of SPLs of Java programs [7]. A prototypical implementation of `DeltaJ` that supports Java 1.5, called `DeltaJ 1.5`, is available [22, 44]. In this section we briefly illustrate the use of `DeltaJ 1.5` in the implementation of a simple product line—the Album PL. Each product in the Album PL is a Java 1.5 program implementing a photo album. This running exam-

```

1  import java.util.Set;
2  import com.finefit.sut.*;
3  import com.finefit.sut.FineFitDriver;
4
5  public class PhotoAlbumDriver extends FineFitDriver {
6      private ArrayPhotoAlbum sut;
7      protected void setup_sut() { sut = new ArrayPhotoAlbum(5); }
8      public void init_sut(com.finefit.model.State args) { }
9      public State retrieve() { return sut.retrieve(); }
10     protected void setup_operation_table() {
11         ops.put("addPhoto", new Operation() {
12             public void apply(com.finefit.model.State args, State outputs) throws Exception {
13                 String id = args.getArg("p");
14                 Photo p = sut.addPhoto(id);
15                 IdMap.instance().associate(p, id);
16             } });
17         ops.put("viewPhotos", new Operation() {
18             public void apply(com.finefit.model.State args, State outputs) throws Exception {
19                 outputs.add_output("result!", 1);
20                 Set<Photo> photos = sut.viewPhotos();
21                 for(Photo p : photos) {
22                     outputs.get_output("result!").add(IdMap.instance().obj2atom(p));
23                 }
24             } });
25     }
26     protected void setup_exception_table() {
27         exceptions.put("PhotoAlbum$PhotoExists", "PHOTO_EXISTS");
28         exceptions.put("PhotoAlbum$AlbumIsFull", "ALBUM_FULL");
29     }
30 }

```

Listing 3: Java code implementing the driver for the Base Album




Fig. 4 Feature model of the Album PL

ple aims at presenting the proposed approach for integrating refinement-based testing into delta-oriented SPLs development rather than at providing a realistic case study.

In delta-oriented programming each product is described by a set of features, where a feature is an abstract description of functionality [4]. Figure 4 depicts the feature model⁶ of the Album PL as a feature diagram. The feature Base is mandatory, while all other features are optional. A product that has the feature Groups requires also the feature Owner. The Album PL has therefore six products. The product that has only the Base feature is the Base Album introduced in Section 2.

A DeltaJ product line consists of a *code base* and a *product-line declaration*. The code base (described in Section 3.1) consists of a set of delta modules describing modification to Java programs. The product-line declaration (described in Section 3.2) provides the connection of the delta modules with the product features; i.e., it specifies which delta modules must be used to generate each product.

⁶ A *feature model* defines the valid feature configurations of an SPL, i.e., the feature configurations that describe the products (see, e.g., [4]).

3.1 DeltaJ delta modules

A DeltaJ delta module describes the changes in a given product (i.e., a Java program) that are needed to implement other products (i.e., other Java programs). The alterations inside a delta module act both at the class/interface level by

- adding or removing a Java *compilation unit*, that is, an interface or a class together with a package declaration and a list of import-statements, or
- modifying the package declarations, or
- modifying the import-statements;

and at the class/interface structure level by

- modifying existing interfaces (i.e., changing the super interfaces and adding, removing or modifying method signatures or nested types), or
- modifying the internal structure of existing classes (i.e., changing the super class or the implemented interfaces and adding, removing or modifying constructors, fields, methods or nested types).

Modifying a method m of a class C means replacing the method body with a new one. The new body may contain the call **original**(\dots), that is replaced in the generated product by a call to a new method with a fresh name m' , which has the same type and body as m before the modification. The new method m' is added to the class C when the product is generated.⁷

3.2 DeltaJ product-line declaration

Listing 4 illustrates the declaration for the Album PL. The product line declaration:

- Lists the product features.
- Lists the delta modules.
- Describes the set of valid feature configurations by means of propositional constraints over the set of features (see, e.g., [4]). For each feature ϕ in the set $\{\overline{\phi}\}$ of the product line features we introduce a propositional variable with the same name. A propositional formula P characterizes a set of feature configurations $\Psi \subseteq \mathcal{P}(\{\overline{\phi}\})$ if and only if, for every feature configuration $\{\overline{\psi}\} \in \Psi$, the formula P is true when the variables in $\{\overline{\psi}\}$ are true and the variables in $\{\overline{\phi}\} \setminus \{\overline{\psi}\}$ are false. The propositional formula in Listing 4 (Base & (Groups \Rightarrow Owner)) represents the six valid feature configurations described by the feature diagram (Fig. 4).
- Describes the relation between delta modules and features by means of a totally ordered set of constraints called *partitions*. A partition consists of a set of *when-clauses*. When-clauses in the same partition are separated by a comma and the end of each partition is indicated by a semicolon. Consider $p \geq 1$ partitions, such that the i -th partition contains $q_i \geq 1$ delta clauses $S_{i,j}$ when $P_{i,j}$ ($1 \leq i \leq p$ and

⁷ This mechanism is similar to the **Super**(\dots) call of FOP [5] and to the *around advice* and *proceed* mechanisms of *aspect-oriented programming* (AOP)—see, e.g., [37,7] for a comparison between DOP and AOP.

```

1  SPL Album {
2    Features = {Base, Remove, Owner, Groups}
3    Deltas = {DBase, DRemove, DOwner, DRemoveAndOwner, DGroups}
4    Constraints { Base & (Groups => Owner); }
5    Partitions {
6      {DBase} when (Base);
7      {DRemove} when (Remove), {DOwner} when (Owner);
8      {DRemoveAndOwner} when (Remove & Owner);
9      {DGroups} when (Groups);
10   }
11   Products {
12     A.Base = {Base};
13     A.Remove = {Base, Remove};
14     A.Owner = {Base, Owner};
15     A.RemoveOwner = {Base, Remove, Owner};
16     A.OwnerGroups = {Base, Owner, Groups};
17     A.OwnerRemoveGroups = {Base, Remove, Owner, Groups};
18   }
19 }

```

Listing 4: DeltaJ 1.5 product-line declaration for the Album PL

$1 \leq j \leq q_i$). Each when-clause consists of a set of delta module names $S_{i,j}$, followed by the **when** keyword and by a propositional formula over features $P_{i,j}$. The sets of delta module names $S_{i,j}$ are pairwise disjoint and their union consists of all the delta modules of the SPL (i.e., the set of sets of delta module names $\{S_{i,j} \mid 1 \leq i \leq p \text{ and } 1 \leq j \leq q_i\}$ represents a partition of the delta module set of the SPL).⁸ The formula $P_{i,j}$ (called the *application condition* of the delta modules in $S_{i,j}$) describes for which feature configurations the delta modules must be applied. Only valid feature configurations (according to the feature model) may be used for product generation—hence the application conditions must be read by assuming that the formula describing the set of valid feature configurations holds. Delta modules in the same partition can be applied in any order whereas the order between partitions must be respected. The ordering allows the designer to enforce semantic requires-relations that are necessary for the applicability of the delta modules.

- Declares some products that can be generated by giving a name to the associated feature configurations (this allows developers to disable the generation of some products without changing the set of valid feature configurations).

According to the Album PL declaration in Listing 4: The features and the valid feature configurations are those described by the feature model in Fig. 4. The delta modules DBase, DRemove, DOwner, and DGroups are associated with the features Base, Remove, Owner, and Groups, respectively; moreover, when the two features Remove and Owner have to be realized, also the delta module DRemoveAndOwner must be applied. All the products corresponding to the six valid feature configurations can be generated.

The delta module DBase (Listing 5) introduces the class Photo, the interface PhotoAlbum, and the class ArrayPhotoAlbum representing the product which has

⁸ In DELTAJ 1.5 each constraint “ $S_{i,1}$ when $P_{i,1}$, ..., S_{i,q_i} when P_{i,q_i} ” is called “partition”, since the set of sets of delta module names $\{S_{i,j} \mid 1 \leq j \leq q_i\}$ is a partition of $\cup_{1 \leq j \leq q_i} S_{i,j}$.

```

1 delta DBase {
2   adds { package it.unito.Album; /* Same as in Listing 1 (top) */ }
3   adds { package it.unito.Album; /* Same as in Listing 1 (middle)*/ }
4   adds { package it.unito.Album; /* Same as in Listing 1 (bottom)*/ }
5 }

```

Listing 5: Code base of the Album PL: delta module DBase

```

1 delta DRemove {
2
3   modifies it.unito.Album.PhotoAlbum {
4     adds public void removePhoto(int location);
5   }
6
7   modifies it.unito.Album.ArrayPhotoAlbum {
8     adds public void removePhoto(int location) {
9       if ((location < 0) || (size <= location))
10        throw new IllegalArgumentException("IllegalLocation");
11       photoAt[location] = photoAt[size-1];
12       photoAt[size-1] = null;
13       size = size -1;
14     }
15   }
16 }
17

```

Listing 6: Code base of the Album PL: delta module DRemove

```

1 delta DOwner {
2
3   adds { package it.unito.Album;
4     class User {
5       ... /* Fields, constructor, methods */
6     }
7   }
8
9   modifies it.unito.Album.PhotoAlbum {
10    adds public void login(String name, String password);
11    adds public void logout();
12    adds nested { public class AlreadyLogged extends RuntimeException {} }
13    adds nested { public class AuthFailed extends RuntimeException {} }
14    adds nested { public class NotAuthorized extends RuntimeException {} }
15    adds nested { public class OwnerNotLoggedIn extends RuntimeException {} }
16  }
17
18  modifies it.unito.Album.ArrayPhotoAlbum {
19
20    ... /* Adds imports, adds fields, modifies constructor, adds methods */
21
22    modifies addPhoto(String image) {
23      if (!isOwnerLoggedIn()) throw new OwnerNotLoggedIn();
24      return original(image);
25    }
26  }
27 }
28

```

Listing 7: Code base of the Album PL: delta module DOwner


```

1  delta DRemoveAndOwner {
2
3      modifies it.unito.Album.ArrayPhotoAlbum {
4          modifies removePhoto(int location) {
5              if (!isOwnerLoggedIn()) throw new OwnerNotLoggedIn();
6              original(location);
7          }
8      }
9
10 }

```

Listing 8: Code base of the Album PL: delta module DRemoveAndOwner

```

1  delta DOwner {
2
3      adds { package it.unito.Album;
4          /* Imports */
5          class Groups {
6              ... /* Fields, constructor, methods */
7          }
8      }
9
10     modifies it.unito.Album.Photo {
11         ... /*Adds a field (group) and two methods (getGroup and setGroup) */
12     }
13
14     modifies it.unito.Album.PhotoAlbum {
15         adds public User updateUser(String name, String password);
16         adds public Group updateGroup(String name, Set<String> memberNames);
17         adds public void removeUser(String name);
18         adds public void removeGroup(String name);
19         adds public void updatePhotoGroup(int location, String groupName);
20         adds nested { public class MissingUser extends RuntimeException {} }
21         adds nested { public class MissingUsers extends RuntimeException {} }
22         adds nested { public class MissingGroup extends RuntimeException {} }
23         adds nested { public class RemoveOwnerGroup extends RuntimeException {} }
24         adds nested { public class RemoveOwner extends RuntimeException {} }
25     }
26
27     modifies it.unito.Album.ArrayPhotoAlbum {
28         ... /* Adds fields, modifies constructor, modifies methods, adds methods */
29     }
30
31 }

```

Listing 9: Code base of the Album PL: delta module DGroups

only the Base feature (cf. Section 2.1). The delta module DRemove (Listing 6) introduces the Remove feature which enables to remove a photo from the album. The delta module DOwner (Listing 7) introduces the Owner feature which protects the photo album by enabling only the owner (who has to login) to modify the album. The delta module DRemoveAndOwner (Listing 8) introduces the code needed to handle the combination of the optional features DRemove and DOwner. The delta module DGroups (Listing 9) introduces a further level of protection by requiring the owner to create users and to associate to each photo the group of users that can view it. The complete code of the delta modules of the Album PL is available at the DeltaFineFit home page [43].

3.3 DeltaJ generation of the products

A product is *valid* if it corresponds to a valid feature configuration. The *product generation mapping* is the mapping that associates each valid feature configuration to the corresponding product (i.e., the Java program obtained by applying the delta modules with a satisfied when-clause to the empty program). This mapping may be partial (since a non-applicable delta module may be encountered during product generation, resulting in an undefined product) and ambiguous (since, for a given feature consideration, two different orders of the delta-modules that are compatible with the order of the partitions may generate two different products).

A delta module is applicable to a Java program if suitable syntactic conditions are satisfied. E.g.:

- each class or interface to be added does not exists;
- each class or interface to be removed or modified exists;
- for every interface to be modified: each method to be added does not exist, each method to be removed exists; and
- for every class to be modified: each method or field to be added does not exist, each method or field to be removed exists, and each method to be modified exists and has the same signature as in the method-modify operation.

A suitable type system could guarantee that if a DeltaJ product line is well-typed, then its product generation mapping is total and unambiguous and all its products are well-typed Java programs. Such a type system has been formalized for the minimal core calculus IMPERATIVE FEATHERWEIGHT DELTA JAVA (IFΔJ) [7]. However, it is not yet implemented in DeltaJ 1.5 [22,44], where the products are generated by considering delta modules in the order in which they occur in the product line declaration,⁹ and most type errors in the products are detected only when the product is generated.

3.4 Delta-oriented SPL development

Delta-oriented programming is a transformational approach for the development of SPLs [39]. For instance, it supports developing an SPL by starting from at least one complete product, called the *core product*, and writing program transformations (the delta modules) that specify changes to be applied to the core product in order to implement other products. In the following we use the phrase *core delta module* to refer to a delta module that when applied to the empty product generates the code of a complete product.

The main advantage of transformational approaches over compositional ones (such as, e.g., FOP [5,3,14,2]) is that when developing an SPL starting from a set of core products, the latter approaches require to begin from simple products (implementing minimal sets of features) to be extended in order to implement the other products. Instead, DOP supports also developing an SPL starting from complex products (implementing arbitrary large sets of features) and transforming them into simpler

⁹ Thus ruling out ambiguity.

```

1  SPL Album {
2      Features = ... /* Same as in Listing 4 */
3      Deltas = ... /* Same as in Listing 4 */
4      Constraints { ... /* Same as in Listing 4 */ }
5      Partitions {
6          {Dall} when (Base);
7          {DnoGroups} when (!Groups);
8          {DnoOwner} when (!Owner);
9          {DnoRemove} when (!Remove);
10     }
11     Products { ... /* Same as in Listing 4 */ }
12 }

```

Listing 10: DeltaJ 1.5 product-line declaration for the complex-core implementation of the Album PL

products by removing features (see [36,38]). More general, DOP is well-suited to support the following SPL development approaches: *proactive* (i.e., all reusable artifacts are planned and developed in advance), *reactive* (i.e., only a basic set of products is planned and developed—when new customer requirements arise, the existing SPL is evolved), and *extractive* (i.e., turning a set of existing applications into an SPL); whereas FOP supports only the proactive approach. Note that the simple-core approach roughly corresponds to the proactive approach, while the complex-core approach may be seen as a particular case of the extractive approach. Proactive SPL development is the most appealing choice from the quality point of view. However, it requires a high upfront investment. Krueger [24] proposed therefore the reactive and extractive approaches to reduce the adoption barrier for SPL engineering.

The implementation of the Album PL illustrated above follows the simple-core approach. Listing 10 illustrates the SPL declaration of an implementation of the Album PL that follows the complex-core approach. The delta modules (their code is available in the DeltaFineFit home page [43]) are such that Dall introduces the code of the product with all the features, while DnoGroups, DnoOwner and DnoRemove remove the code that implements the features Groups, Owner and Remove, respectively.

4 Delta-oriented specification of SPLs

In this section we describe *delta tables*, a concept that we have developed and used to derive tabular specifications of products by specifying the difference between the product specifications.

4.1 Delta tables and delta-table modules

In Section 2.2 we have described the tables used by the FineFit model. We refer to them hereafter as *ordinary tables*. Figure 2 presents the structure of each ordinary table. The idea of delta tables is to derive a new table from an existing one, based on the difference (the delta) between the tables. This difference is provided in an easily readable way.

A delta table is recognized by its name (the upper-left cell), which is preceded with the Δ symbol. Delta tables have the same structure as the ordinary tables, yet their cells may contain *delta operators* (namely, match “*”, remove “—”, insert “►”, and replace “*►”). Let us denote the set of ordinary tables by T_O and the set of delta tables by T_Δ . A delta table t_Δ can be *applied* to an ordinary table t_o , written as $apply(t_o, t_\Delta)$, resulting in a new ordinary table. That is, $apply$ is a function $apply: T_O \times T_\Delta \rightarrow T_O \cup \varepsilon$, where ε denotes an empty (non-existent) table in the case that the table is removed by the application. Following is an example of a delta table application which yields a resulting table (to be explained later). We use the infix notation: $T_O \text{ apply } T_\Delta$.

f(x)	$x < 0$	$x \geq 0$
y	x	x
z	1	0

 $apply$

$\Delta f(x)$	$x < 0$	$x \geq 0$
y	*►-x	*
—	—	—

yields

f(x)	$x < 0$	$x \geq 0$
y	-x	x

A *delta-table module* (DTM) is a set of delta and non-delta tables. Delta table modules specify the changes to the FineFit specification of a product that are needed to obtain the specification of another product (i.e., another set of tables). For this we introduce the function $Apply: \mathcal{P}(T_O) \times \mathcal{P}((T_O \cup T_\Delta)) \rightarrow \mathcal{P}(T_O)$, where $\mathcal{P}(T_O)$ and $\mathcal{P}(T_\Delta)$ denote sets of tables.

To define the $Apply$ function we must distinguish between a table and its name, since when overriding a table a with a table b the tables may be different, i.e. $a \neq b$, yet they must share the same name, e.g., specify the same operation. Let \bar{t} denote the name of table t and \bar{T} denote the set of all table names of the set of tables T . A delta table module $B \in \mathcal{P}((T_O \cup T_\Delta))$ is applied to a FineFit specification $A \in \mathcal{P}(T_O)$ as follows:

$$Apply(A, B) = \bigcup_{a \in A, b \in B} \begin{cases} \{a\} & \text{if } \bar{a} \notin \bar{B} \\ \{b\} & \text{if } b \in T_O \\ \{apply(a, b)\} & \text{if } b \in T_\Delta, \Delta \bar{a} = \bar{b} \text{ and } apply(a, b) \neq \varepsilon \end{cases}$$

The three different cases are: (case 1) tables in original table set that are not present in the DTM are copied to the resulting set of tables; (case 2)¹⁰ ordinary tables of the DTM are copied to the resulting table (and may overwrite tables from the original set A); and (case 3) if a table is marked with the Δ symbol, then it is a delta table and it is applied to the corresponding table in the original set. When $apply(a, b)$ returns ε , a is effectively removed. Note that multiple conditions may be satisfied simultaneously; in this case all tables whose conditions are satisfied are included.

The first version of delta tables, as proposed in [10], is capable of describing in a concise and an intuitive way the modification of value and variable cells and the refinement of condition cells. The latter means that subconditions can be introduced below existing condition cells of a table. We realized, however, that this condition refinement is not sufficient for handling all modifications that are necessary in practice, as it is often required to change the condition hierarchy in a more complex way.

¹⁰ Note that tables satisfying this case do not satisfy case 1.

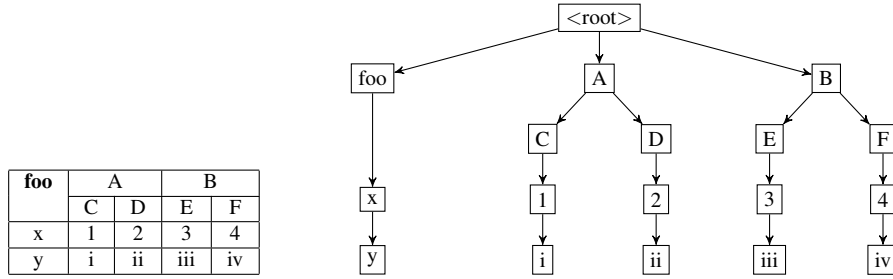


Fig. 5 Table foo and its tree structure representation (used for explanations)

We have investigated several syntactic notations; for example, dividing the condition cells into sections with different semantics of applications. Those notations, however, either have been complex or their application process was hard to understand.

The challenge of defining transformations on Parnas tables results from the fact that Parnas tables have both a tabular and a hierarchical structure. When defining operations, we have to decide whether the tables should be treated as matrices or as trees. We found that embedding a matrix in a tree is easier than the other way around, and it allows us to define basic operations that are applicable to all kinds of tables and to all kinds of cells. In particular, these operations allow a concise treatment of hierarchical conditional cells which is the most difficult feature to design. These operations are defined by the *apply* function that is introduced in Section 4.2. To handle also the matrix structure of the tables in a concise way and to provide also abbreviations and convenience rules for the user, we define another function—“prepare”, which is described in Section 4.3. A strength of the approach is that it has one core algorithm that is relatively small and is recursively and uniformly applied on all kinds of cells of the table. Furthermore, the same algorithm is applied to all kinds of tables in our framework.

4.2 Hierarchical representation of tables and operations of delta tables

Delta tables have the same syntax as the tables shown in Fig. 2. However, their cells may contain special symbols representing operations that are executed by the *apply* function. The operations and their corresponding symbols are: match “*”, remove “—”, insert “►”, and replace “*►”. In order to deal with the condition hierarchy of operation tables, the *apply* function treats tables as ordered trees rather than as matrices, and traverses them recursively from top to bottom. Consider, for example, the table foo and its encoding as a tree (Fig. 5). The columns of the table are the branches of the tree. A table cell c_1 which is located directly below cell c_0 is a child of c_0 in the tree representation. If c_0 spans several cells, e.g., c_1, \dots, c_k , those cells are children of c_0 . The order of siblings must be the same as in the table. Implicitly there also is present a $\langle \text{root} \rangle$ cell, whose children are the top level cells of the table. Note that each cell represents a tree, namely the subtree which has the cell as its root. Hence, the entire tree is represented by the $\langle \text{root} \rangle$ cell.

In the following we define the *apply* function and respectively the semantics of the delta table operations which are interpreted by this function. We use the infix notation of *apply* and write \boxed{A} to represent a cell with content A. The notation

$$c_0: \begin{array}{|c|c|c|} \hline C \\ \hline c_1 & \dots & c_k \\ \hline \end{array}$$

represents a node c_0 with value C and with subnodes c_1, \dots, c_k .¹¹ The values of c_i are denoted by $op(c_i)$. We omit writing “ $c_0 :$ ” when possible. Given a table t_o and a delta table t_Δ , $t_o \text{ apply } t_\Delta$ yields a resulting table t_r by simultaneously traversing the structures of t_o and t_Δ .

In Appendix B we provide the algorithm of the *apply* function and additional details. Here we present a more intuitive definition, using graphical notation and examples.

4.2.1 Definition of *apply* : $T_O \times T_\Delta \rightarrow T_O$

For each table $t_o \in T_O$ with content A and subnodes o_1, \dots, o_m , and for each $t_\Delta \in T_\Delta$ with subnodes d_1, \dots, d_n and operation $\delta = op(t_\Delta)$ where $\delta \in \{*, -, \blacktriangleright, * \blacktriangleright\}$, $t_o \text{ apply } t_\Delta$ is defined inductively over the tree structure representation of the tables.

- Case “*”: The match operator acts as a placeholder that matches any cell and copies it to the resulting table.

$$o_0: \begin{array}{|c|c|c|} \hline A \\ \hline o_1 & \dots & o_m \\ \hline \end{array} \quad \text{apply} \quad \begin{array}{|c|c|c|} \hline * \\ \hline d_1 & \dots & d_n \\ \hline \end{array} \quad \text{yields} \quad \begin{array}{|c|c|c|} \hline A \\ \hline r_1 & \dots & r_n \\ \hline \end{array}$$

The branches r_1, \dots, r_n of the resulting table are defined as:

$$r_i := \text{if } op(d_i) \neq \blacktriangleright \text{ then } apply(o_j, d_i) \text{ else } apply(o_0, d_i), \quad (1)$$

where \blacktriangleright is the insert operator, $j = \min(i, m)$, and $op(d_i)$ denotes the delta table operator of node d_i .

Remark 1 The same definition of r_i applies to all the other cases below. Formally the definition of r_i is the inductive step of the definition of *apply* but it is presented here in the interest of readability.

Remark 2 When the leaf of the condition structure in the delta table is a match operator and the original table’s hierarchy however has further subconditions at this position, then the match operator is applied also to the subconditions of the original table. For example, if the delta table has only one condition cell and it is marked with “*”, then the entire condition hierarchy of the original table is copied to the resulting table.

The placeholder can also be used as a subexpression of a bigger expression, but then it must be embedded in braces “(*)”, as in the following example:

$$\boxed{x} \quad \text{apply} \quad \boxed{(*) + y} \quad \text{yields} \quad \boxed{x + y}$$

¹¹ The bottom line of the cells is removed to indicate that c_i represents a subtree rather than the content of a cell.

- Case “–”: The remove operator matches any cell content A and removes the cell from the resulting table.

A			<i>apply</i>	—			<i>yields</i>	r_1 ... r_n		
o_1	...	o_m		d_1	...	d_n				

The following example combines operations. The first operation removes a cell (i.e., it is not copied to the resulting table) whereas the second copies the cell into the resulting table.

<table><tr><td>A</td></tr><tr><td>D</td></tr></table>	A	D	<i>apply</i>	<table><tr><td>—</td></tr><tr><td>*</td></tr></table>	—	*	yields	<table><tr><td>D</td></tr></table>	D
A									
D									
—									
*									
D									

- Case “►”: The insert operator inserts a cell c_{ins} at the position of the current cell c_i of the original table. The *apply* function proceeds on subsequent cells as if c_i and all its siblings, $c_1, \dots, c_i, \dots, c_n$, are children of c_{ins} . The insert operation is defined partially by the following schema and partially by the recursive step described above (Equation 1).

A			<i>apply</i>	▶I			<i>yields</i>	I		
o_1	...	o_m		d_1	...	d_n		r_1	...	r_n

The following example presents the case where c_{ins} has a subsequent cell:

<table><tr><td>A</td></tr></table>	A	<i>apply</i>	<table><tr><td>►I</td></tr><tr><td>*</td></tr></table>	►I	*	yields	<table><tr><td>I</td></tr><tr><td>A</td></tr></table>	I	A
A									
►I									
*									
I									
A									

We now demonstrate the case where c_{ins} has several subsequent cells:

<table><tr><td>A_1</td><td>\dots</td><td>A_n</td></tr></table>	A_1	\dots	A_n	<i>apply</i>	<table><tr><td colspan="3">\blacktrianglerightI</td></tr><tr><td>*</td><td>\dots</td><td>*</td></tr></table>	\blacktriangleright I			*	\dots	*	yields	<table><tr><td colspan="3">I</td></tr><tr><td>A_1</td><td>\dots</td><td>A_n</td></tr></table>	I			A_1	\dots	A_n
A_1	\dots	A_n																	
\blacktriangleright I																			
*	\dots	*																	
I																			
A_1	\dots	A_n																	

Multiple insert operations can be applied to the current cell of the original table. This results in several subbranches such that the children c_1, \dots, c_n of c_0 occur on each subbranch. This is the desired behavior for refining conditions. For example,

<table><tr><td colspan="2">A</td></tr><tr><td>C</td><td>D</td></tr></table>	A		C	D	<i>apply</i>	<table><tr><td colspan="4">*</td></tr><tr><td colspan="2">►I</td><td colspan="2">►J</td></tr><tr><td>*</td><td>*</td><td>*</td><td>*</td></tr></table>	*				►I		►J		*	*	*	*	<i>yields</i>	<table><tr><td colspan="4">A</td></tr><tr><td colspan="2">I</td><td colspan="2">J</td></tr><tr><td>C</td><td>D</td><td>C</td><td>D</td></tr></table>	A				I		J		C	D	C	D
A																																
C	D																															
*																																
►I		►J																														
*	*	*	*																													
A																																
I		J																														
C	D	C	D																													

- Case “*►”: The replace operation replaces the current cell of the original table with the cell of the delta table. This operation can be simulated by a combination of remove and insert operations; introducing it, however, greatly improves readability. For example,

A			<i>apply</i>	$\ast \blacktriangleright X$			<i>yields</i>	X		
o_1	...	o_m		d_1	...	d_n		r_1	...	r_n

If the replace operation cannot be applied because the original table does not contain a cell at the current position (e.g., it is at the end of the branch), the operation is changed to an insert operation.

Let us demonstrate the application of a delta table that consists of one column and contains all the described operators. The *apply* function simultaneously traverses both the original and the delta tables from top to bottom and applies the delta table operators. For convenience, we write the recursion step numbers of the *apply* function left to the cells.

1:	A		1:	*		1:	A
2,3:	C		2:	►I		2:	I
4:	I	<i>apply</i>	3:	*	<i>yields</i>	3:	C
5:	i		4:	*►D		4:	D
6:	w		5:	—		5:	w
			6:	*		6:	

Here we show the application of a hierarchical delta table with several columns (branches):

A							A		
C	D	<i>apply</i>	►I		►J		<i>yields</i>	I	J
1	2		—	*	*	—		D	C
i	ii		—	*	*►X	—		2	X
			—	*	*►Y	—		ii	Y

Note that the *apply* function treats columns independently and does not prevent vertical misalignment. Delta tables have to be written in such way that the desired outcome is generated, e.g. Parnas tables. The next section describes techniques that help writing delta tables.

4.3 Preprocessing rules

The *apply* function defined in the previous section takes care of the hierarchical nature of Parnas tables and provides a simple core transformation that is uniformly applied to all cells. Here we introduce the *prepare* function, which takes care of the *tabular*¹² nature of Parnas tables and defines a set of rules that simplify writing and reading of delta tables. The *prepare* function acts as a transformation layer on top of the *apply* function, i.e., it compiles a delta table with abbreviations and higher-level constructs into a delta table that uses only the basic operations.

The preprocessing rules are defined by the function $prepare : T_O \times T_\Delta \rightarrow T_\Delta$ which takes as parameters an ordinary table and a delta table and returns a resulting delta table. The resulting delta table is used as input to the *apply* function, as has been defined in the previous section. The function $apply' : T_O \times T_\Delta \rightarrow T_O \cup \varepsilon$ is the composition of the *apply* : $T_O \times T_\Delta \rightarrow T_O \cup \varepsilon$ and the *prepare* functions and is defined as:

$$apply'(t_o, t_\Delta) = apply(t_o, prepare(t_o, t_\Delta)) \quad (2)$$

for all $t_o \in T_O$ and $t_\Delta \in T_\Delta$. The function $Apply' : \mathcal{P}(T_O) \times \mathcal{P}((T_O \cup T_\Delta)) \rightarrow \mathcal{P}(T_O)$ which applies a delta-table module on a set of ordinary tables is similar to the function *Apply* that is described in Section 4.1, but uses the function *apply'* rather than *apply*.

In the following we informally describe the *prepare* function with a set of rules that must be applied in the given order. Some of the rules do not depend on a particular ordinary table, and in this case we omit the ordinary table (i.e., we omit writing the first argument).

¹² Unlike the *apply* function, the *prepare* function is aware of the row-alignment of cells in a table and provides special treatment for different kinds of cells.

Rule 1. When the delta table contains a cell with no associated operation, \blacktriangleright is the default operation. Thus in principle the user does not have to use the \blacktriangleright symbol. However, the user should write \blacktriangleright explicitly when a cell's content is changed to indicate that it is not just overwritten with the same content.

Recall that \blacktriangleright can behave like an insert operation (from the definition of *apply* in Section 4.2.1). Hence, depending on the position of \blacktriangleright and its content it can replace a cell, keep a cell (by replacing it with the same content), or insert a cell. The leading Δ symbol is removed from the name of the delta table to match the name with the original table.

Rule 2. If *all* cells below a condition cell or a name cell c contain the remove operation, c is removed as well. For example:¹³

ΔT	A			
	B		C	
	D	—	F	G
x	—	2	—	—
y	—	ii	—	—

prepare yields

$\blacktriangleright T$	$\blacktriangleright A$			
	$\blacktriangleright B$		—	
	—	—	—	—
$\blacktriangleright x$	—	$\blacktriangleright 2$	—	—
$\blacktriangleright y$	—	$\blacktriangleright ii$	—	—

The first rule is applied here and is responsible for adding the \blacktriangleright operations. The cells A and B are not removed because cells exist below them which are not removed, namely 2 and ii .

This rule enables the user to remove a column of value cells and the corresponding conditions, without having to write “—” in the condition cells. This makes it easier to see which condition is removed. Furthermore, this rule enables to remove an entire table T from the DTM by writing:

ΔT
—

In order to remove all cells of the table except for the name cell, the user may write $\boxed{\Delta T}$.

Rule 3. A variable cell may contain a set of variables. These are expanded by the preprocessing step into a set of rows, one for each variable. For instance:

Given	<table><tr><td>ΔT</td><td>A</td><td>B</td></tr><tr><td>$\{x,y\}$</td><td>1</td><td>3</td></tr></table>	ΔT	A	B	$\{x,y\}$	1	3	, <i>prepare yields</i>	<table><tr><td>$\blacktriangleright T$</td><td>$\blacktriangleright A$</td><td>$\blacktriangleright B$</td></tr><tr><td>$\blacktriangleright x$</td><td>$\blacktriangleright 1$</td><td>$\blacktriangleright 3$</td></tr><tr><td>$\blacktriangleright y$</td><td>$\blacktriangleright 1$</td><td>$\blacktriangleright 3$</td></tr></table>	$\blacktriangleright T$	$\blacktriangleright A$	$\blacktriangleright B$	$\blacktriangleright x$	$\blacktriangleright 1$	$\blacktriangleright 3$	$\blacktriangleright y$	$\blacktriangleright 1$	$\blacktriangleright 3$
	ΔT	A	B															
$\{x,y\}$	1	3																
$\blacktriangleright T$	$\blacktriangleright A$	$\blacktriangleright B$																
$\blacktriangleright x$	$\blacktriangleright 1$	$\blacktriangleright 3$																
$\blacktriangleright y$	$\blacktriangleright 1$	$\blacktriangleright 3$																

Rule 4. The user may write $*$ in a variable cell in order to refer to all variables of the original table that are not explicitly mentioned in the delta table. A row that contains the $*$ operator in its variable cell is called a *default row*. Only one default row may be specified. A default row is added for each variable not mentioned in the variable cells of the delta table.¹⁴ For instance, given

<table><tr><td>T</td><td>A</td><td>B</td></tr><tr><td>x</td><td>1</td><td>3</td></tr><tr><td>y</td><td>2</td><td>4</td></tr></table>	T	A	B	x	1	3	y	2	4	and	<table><tr><td>ΔT</td><td>A</td><td>B</td></tr><tr><td>*</td><td>—</td><td>*</td></tr><tr><td>z</td><td>—</td><td>6</td></tr></table>	ΔT	A	B	*	—	*	z	—	6	<i>prepare yields</i>	<table><tr><td>$\blacktriangleright T$</td><td>—</td><td>$\blacktriangleright B$</td></tr><tr><td>*</td><td>—</td><td>*</td></tr><tr><td>*</td><td>—</td><td>*</td></tr><tr><td>$\blacktriangleright z$</td><td>—</td><td>$\blacktriangleright 6$</td></tr></table>	$\blacktriangleright T$	—	$\blacktriangleright B$	*	—	*	*	—	*	$\blacktriangleright z$	—	$\blacktriangleright 6$
T	A	B																																
x	1	3																																
y	2	4																																
ΔT	A	B																																
*	—	*																																
z	—	6																																
$\blacktriangleright T$	—	$\blacktriangleright B$																																
*	—	*																																
*	—	*																																
$\blacktriangleright z$	—	$\blacktriangleright 6$																																

¹³ Recall that the row span of the name cell defines that the first three rows of the example are condition rows and the other rows contain variable and value cells.

¹⁴ This includes variables checked by the conditional operator.

The actual delta table that is applied in the following example is, hence, the one above, as created by the preprocessing. Recall that according to the definition of *apply*, the replace operations in the last row of this example behave like insert operations.

T	A	B		ΔT	A	B		T	B
x	1	3		*	—	*		x	3
y	2	4		z	—	6		y	4
								z	6

Rule 5. When no default row is specified in the delta table, rows of the original tables with variable names that do not occur in the delta table are copied to the resulting table. This is achieved by filling the cells of the row with the match (*) operation, except for those cells corresponding to columns that are removed by the delta table. A column of a delta table is removed if all its value or variable cells contain the remove operation. For instance, given

T	A	B		ΔT	A	B		* \rightarrow T	—	* \rightarrow B
x	1	3		z	—	6		*	—	*
y	2	4						*	—	*
								* \rightarrow z	—	* \rightarrow 6

Hence:

T	A	B		ΔT	A	B		T	B
x	1	3		z	—	6		x	3
y	2	4						y	4
								z	6

Rule 6. Here we introduce the conditional remove operator $X:-$ which may be used only in the left most column of a table. The operator applies a remove operation to the first occurrence of a cell with content X . If such a cell does not exist in the original table, the entire row of the delta table with the $X:-$ operator is ignored. For example:

Given	A					*
	B					*
	C		C:-			—
	D		*			*
	E					*

The * operator below $C:-$ works as a default row (see Rule 4) and is needed to keep all the other cells. Without the * operator, rule 2 treats $C:-$ simply as a remove operation and the entire column would be removed in this example.

Hence:

A						A
B						B
C		C:-				D
D		*				E
E						

The purpose of the conditional remove operator is to provide a convenient notation to remove cells from an enumeration table.

Rule 7. Rows of the delta table with variable and value cells are sorted such that the variables names match those of the original table.¹⁵ Thus the user does not have to pay attention to the exact order of variables but only to their names. Rows of the delta table with additional variables (that do not exist in the original table) are appended at the end. For example:

Given	<table><tr><td>T</td><td>A</td></tr><tr><td>x</td><td>1</td></tr><tr><td>y</td><td>2</td></tr></table>	T	A	x	1	y	2	and	<table><tr><td>ΔT</td><td>A</td></tr><tr><td>y</td><td>3</td></tr><tr><td>z</td><td>4</td></tr><tr><td>x</td><td>5</td></tr></table>	ΔT	A	y	3	z	4	x	5	, prepare yields	<table><tr><td>$\blacktriangleright T$</td><td>$\blacktriangleright A$</td></tr><tr><td>$\blacktriangleright x$</td><td>$\blacktriangleright 5$</td></tr><tr><td>$\blacktriangleright y$</td><td>$\blacktriangleright 3$</td></tr><tr><td>$\blacktriangleright z$</td><td>$\blacktriangleright 4$</td></tr></table>	$\blacktriangleright T$	$\blacktriangleright A$	$\blacktriangleright x$	$\blacktriangleright 5$	$\blacktriangleright y$	$\blacktriangleright 3$	$\blacktriangleright z$	$\blacktriangleright 4$
	T	A																									
	x	1																									
y	2																										
ΔT	A																										
y	3																										
z	4																										
x	5																										
$\blacktriangleright T$	$\blacktriangleright A$																										
$\blacktriangleright x$	$\blacktriangleright 5$																										
$\blacktriangleright y$	$\blacktriangleright 3$																										
$\blacktriangleright z$	$\blacktriangleright 4$																										

Note that the \blacktriangleright operation is used as the default one due to rule 1. This has the effect that the value cells of corresponding rows are overwritten with values from the delta table, and rows defining new variables are added to the resulting table. For example:

T	A	B
x	1	3
y	2	4

ΔT

A

B

y

5

*

z

6

7

x

*

8

yields

T

A

B

x

1

8

y

5

4

z

6

7

Rule 8. When all value cells in a row are removed, also the variable cell of that row is removed. This provides an intuitive way to remove variables. For instance:

Given	ΔT	A	B
	x	—	—
	y	5	6

, prepare yields

$\blacktriangleright T$	$\blacktriangleright A$	$\blacktriangleright B$
—	—	—
$\blacktriangleright y$	$\blacktriangleright 5$	$\blacktriangleright 6$

Hence:

T	A	B
x	1	3
y	2	4

apply'

ΔT	A	B
x	—	—
y	5	6

yields

T	A	B
y	5	6

Rule 9. The insert operation has been defined in a way that enables refinement of conditions (see Section 4.2). Recall that if a cell c_{ins} is inserted at the subtree where the corresponding original subtree has cells c_1, \dots, c_n , then c_1, \dots, c_n are treated as children of c_{ins} for subsequent operations. For example:

<table><tr><td>A_1</td><td>A_2</td></tr></table>	A_1	A_2	$apply'$	<table><tr><td colspan="2">$\blacktriangleright I$</td><td colspan="2">$\blacktriangleright J$</td></tr><tr><td>*</td><td>*</td><td>*</td><td>*</td></tr></table>	$\blacktriangleright I$		$\blacktriangleright J$		*	*	*	*	yields	<table><tr><td colspan="2">I</td><td colspan="2">J</td></tr><tr><td>A_1</td><td>A_2</td><td>A_1</td><td>A_2</td></tr></table>	I		J		A_1	A_2	A_1	A_2
A_1	A_2																					
$\blacktriangleright I$		$\blacktriangleright J$																				
*	*	*	*																			
I		J																				
A_1	A_2	A_1	A_2																			

The above example is problematic if A_1 is a name cell, i.e., it contains the table's name, and I and J are condition cells. In this case, the name cell A_1 becomes a child of the condition cells I and J which clearly is wrong.

foo()	true
x	y

apply'

$\blacktriangleright y < 0$	$\blacktriangleright y \geq 0$
*	*
*	*

yields

$\blacktriangleright y < 0$		$\blacktriangleright y \geq 0$	
foo()	true	foo()	true
x	y	x	y

¹⁵ Rows with conditional remove operators are also moved to the right place by this rule, i.e., where the content of a cell matches the condition.

Hence, in order to obtain a well-formed operation table (Section 2.2.2), the first column below the inserted condition cell must be removed when refining the top-level conditions of a table. To release the user from this task and to improve readability of the tables, the following preprocessing rule is applied. If c is a top-level table cell, i.e. a child of the root node, containing an insert operation (say $\blacktriangleright C$), then given:

$$c: \begin{array}{|c|c|c|} \hline & \blacktriangleright C & \\ \hline op_1 & \dots & op_n \\ \hline \end{array} \quad \text{prepare yields} \quad c: \begin{array}{|c|c|c|} \hline & \blacktriangleright C & \\ \hline - & op_1 & \dots & op_n \\ \hline \end{array}$$

where $op_1 \dots op_n$ are delta operations. Once this rule is used, the delta application in the following example is correct:

$$\begin{array}{|c|c|} \hline foo() & true \\ \hline x & y \\ \hline \end{array} \quad \text{apply}' \quad \begin{array}{|c|c|c|} \hline foo() & \blacktriangleright y < 0 & \blacktriangleright y \geq 0 \\ \hline & * & * \\ \hline x & -y & * \\ \hline \end{array} \\ \text{yields} \\ \begin{array}{|c|c|c|} \hline foo() & y < 0 & y \geq 0 \\ \hline & true & true \\ \hline x & -y & y \\ \hline \end{array}$$

The “true” has not been removed for demonstration purpose, though it easily can be removed using the remove operation above or below the insert operation. Another possibility to remove the cells with the “true” condition is to use the remove operation instead of the insert operation. This works despite the fact that the second remove operation does not have a corresponding cell in the original table.

5 Specifying and testing the Album PL with DeltaFineFit

In this section we outline how to specify and test a delta-oriented SPL using DeltaFineFit. First (in Section 5.1) we illustrate how to write the DeltaFineFit specification. Then (in Section 5.2) we show how to extend a DeltaJ SPL in order to generate for each product a version of the product equipped with the code for using FineFit. Finally (in Section 5.3) we outline how DeltaFineFit can be used effectively to test an SPL.

5.1 DeltaFineFit specification

A DeltaFineFit specification of an SPL is a product line where each product is a FineFit specification of a Java program (cf. Section 2.2). It consists of a product line declaration and a set of delta-table modules. Delta-table modules are sets of delta and non-delta tables (cf. Section 4.1), and a DeltaFineFit product-line declaration is a DeltaJ product line declaration where delta-table modules are used instead of delta-modules. Notably, a DeltaJ product-line declaration can be understood as a DeltaFineFit product-line declaration by providing a one-to-one mapping between delta-table modules and delta modules.

The DeltaFineFit generation procedure is analogous to the DeltaJ 1.5 generation procedure of a Java program (outlined in Section 3.3): given a valid feature configuration the DTMs with a satisfied when-clause are applied to the empty FineFit specification in the order in which they appear in the DeltaFineFit product-line declaration. This procedure describes a mapping that associates each valid feature configuration to the corresponding FineFit specification, and this mapping may be partial.

The tables describing the model of the Base Album (see Fig. 3) can be seen as the DTMBase delta-table module that when applied to the empty abstract model produces the model of the Base Album. The other DTMs of the Album PL specification are listed in Figures 6, 7, 8, and 9. The delta-table module DTMRMRemove (Fig. 6) specifies the Remove feature which enables to remove a photo from the album. The delta-table module DTMOwner (Fig. 7) specifies the Owner feature which protects the photo album by enabling only the owner (who has to login) to modify the album. The DTMRMRemoveAndOwner delta-table module (Fig. 8) introduces the needed modifications to specify the combination of the two optional features Remove and Owner. The delta-table module DTMGroups (Fig. 9) specifies a further level of protection by requiring the owner to create users and to associate each photo with the group of users that can view it.

The simple-core implementation of the Album PL (Listing 4) can therefore be understood as the DeltaFineFit product-line declaration for the specification of the Album PL, modulo the mapping that associates the delta-table modules DTMBase, DTMRMRemove, DTMOwner, DTMRMRemoveAndOwner, and DTMGroups to the delta modules DBase, DRemove, DOwner, DRemoveAndOwner, and DGroups, respectively.

The specification of the Album PL (Figures 6 – 9) uses the simple-core approach (cf. Section 3.4). The declaration of the complex-core DeltaJ 1.5 implementation of the Album PL given in Listing 10 can be used as a declaration of a complex-core DeltaFineFit specification of the Album PL. This is achieved by associating to the delta modules Dall, DnoGroups, DnoOwner, and DnoRemove the appropriate delta-table modules DTMall, DTMnoGroups, DTMnoOwner, and DnoRemove, respectively. The tables of the complex-core specification are available at the DeltaFineFit home page [43]. Note that the tables of the core DTM for the product with all features, DTMall, are the tables in the specification of the product with all features.

In the following examples we take a close look at some of the specification tables from the figures.

Example 1 This example shows the application of the Δ **addPhoto** delta-operation table of the delta-table module DTMOwner to the operation table **addPhoto** of the base product. The original operation table (from Fig. 3) is defined as:

addPhoto(p?:Photo, report!:Report)	#photoAt < MAX		#photoAt ≥ MAX
	p? !in ran[photoAt]	p? in ran[photoAt]	true
photoAt	photoAt.add[p?]	=	=
report!	OK	PHOTO_EXISTS	ALBUM_FULL

The delta-operation table (at the bottom of Fig. 7) is defined as:

Δ addPhoto(p?:Photo, report!:Report)	► ownerName in loggedIn		► ownerName !in loggedIn	
	*	*	-	true
	*	*	-	true
report!	*	*	-	AUTH.FAILED

The application results in the following operation table:

addPhoto(p?:Photo, report!:Report)	ownerName in loggedIn		ownerName !in loggedIn	
	#photoAt < MAX	#photoAt ≥ MAX	true	
	p? !in ran[photoAt]	p? in ran[photoAt]	true	true
photoAt	photoAt.add[p?]	=	=	=
report!	OK	PHOTO.EXISTS	ALBUM.FULL	AUTH.FAILED

Then, applying to the above table the following delta-operation table Δ addPhoto of the delta-table module DTMGrouPs (from Fig. 9)

Δ addPhoto(p?:Photo, report!:Report)	*			
	*			*
	*			*
groupPhotos	groupPhotos + p? → OWNER_GROUP	=	=	=

yields an operation table with an additional row (which belongs to the specification of the product with features Base, Owner and Groups). Here we use a new syntax, where condition cells with a match operation span multiple cells of a column. These condition cells in the corresponding column of the original table are copied to the resulting table.¹⁶

Note that using the following table to add the additional row is not correct, because multiple variable cells occur below a condition cell:

Δ addPhoto(p?:Photo, report!:Report)	*			
groupPhotos	groupPhotos + p? → OWNER_GROUP	=	=	=

This is not a valid operation table and delta tables must satisfy the syntax of ordinary tables as defined in Fig. 2.

Example 2 Next, we consider applying the delta-operation table Δ login of DTMGrouPs to the operation table login of DTMOwner. The original operation table login (last but third table in Fig. 7) is:

login(n?:Name, p?:Password, report!:Report)	no loggedIn		some loggedIn
	n? = ownerName and p? = ownerPassword	n? != ownerName or p? != ownerPassword	true
loggedIn	loggedIn + n?	=	=
report!	OK	AUTH.FAILED	ALREADY.IN

The delta-operation table Δ login (third table in Fig. 9) is defined as:

Δ login(n?:Name, p?:Password, report!:Report)	*		*
	*► n? in dom[users] and passwords[users[n?]] = p?	*► n? !in dom[users] or passwords[users[n?]] != p?	*
	► p? = ownerPassword	► true	► true
loggedIn	loggedIn + users[n?]	=	=

removePhoto(l?:Int, report!::Report)	l? in dom[photoAt]	l? !in dom[photoAt]	Δ Report
photoAt	photoAt.delete[l?]	=	NO_PHOTO
report!	OK	NO_PHOTO	

Fig. 6 Delta-table module DTMRemove: it enables to remove a photo from the album

ΔAtom	Scope	Δ Report	ΔState variable	Type
Name	4	AUTH_FAILED	ownerName	Name
Password	4	ALREADY_IN	ownerPassword	Password
			loggedIn	set Name

ΔInvariant name	Invariant
AtMostOneLoggedIn	lone loggedIn

Δ init(n?:Name, p?:Password)	true
ownerName	n?
ownerPassword	p?
loggedIn	none

login(n?:Name, p?:Password, report!::Report)	no loggedIn		some loggedIn
	n? = ownerName and p? = ownerPassword	n? != ownerName or p? != ownerPassword	true
loggedIn	loggedIn + n?	=	=
report!	OK	AUTH_FAILED	ALREADY_IN

logout()	true
loggedIn	none

Δ addPhoto(p?:Photo, report!::Report)	► ownerName in loggedIn		► ownerName !in loggedIn	
	*	*	-	true
	*	*	-	true
report!	*	*	-	AUTH_FAILED

Fig. 7 Delta-table module DTMOwner (it contains a delta-atoms table, a delta-elements table, a delta-state table, two operation tables, and one delta-operation table): only the owner can add a photo

Δ removePhoto(l?:Int, report!::Report)	► ownerName in loggedIn		► ownerName !in loggedIn	
	*	*	-	true
report!	*	*	-	AUTH_FAILED

Fig. 8 Delta-table module DTMRemoveAndOwner: a photo can be removed only by its owner

Δ Atom	Scope
User	4
Group	4

Δ State variable	Type
groups	Name \rightarrow Group
users	Name \rightarrow User
members	Group \rightarrow User
passwords	User \rightarrow Password
groupPhotos	Photo \rightarrow Group
ownerPassword	Name

Δ Report
WRONG_PASSWORD
ALREADY_IN
MISSING_USERS
MISSING_USER
REM_OWNER
REM_OWNER_GROUP
NO_GROUP
NO_PHOTO

Δ Invariant name	Invariant
validMembers	members in ran[groups] \rightarrow ran[users]
validPasswords	passwords in ran[users] \rightarrow one Password
validGroups	groupPhotos in Photo \rightarrow ran[groups]

Δ init(o?:User, p?:Password, n?:Name, ogn?:Name, og?:Group)	true
ownerPassword	—
passwords	$o? \rightarrow p?$
users	$n? \rightarrow o?$
groups	$ogn? \rightarrow og?$
members	$og? \rightarrow o?$
ownerGroupName	$ogn?$

Δ addPhoto(p?:Photo, report!:Report)	*			
	*	*	*	*
groupPhotos	groupPhotos + p? \rightarrow OWNER_GROUP	=	=	=

Δ viewPhotos(result!: seq Photo, report!:Report)	-	
	one loggedIn	no loggedIn
result!	photoAt.elems	none
report!	OK	AUTH_FAILED

Δ login(n?:Name, p?:Password, report!:Report)	*		*
	$* \blacktriangleright n? \text{ in dom[users] and passwords[users[n?]] = p?}$	$* \blacktriangleright n? \text{ !in dom[users] or passwords[users[n?]] != p?}$	*
	$\blacktriangleright p? = \text{ownerPassword}$	$\blacktriangleright \text{true}$	$\blacktriangleright \text{true}$
loggedIn	loggedIn + users[n?]	=	=

The operation tables **updateUser**, **updateGroup**, **removeUser**, **removeGroup** and **updatePhotoGroup** are available at the DeltaFineFit home page [43]

Fig. 9 Delta-table module DTMGroups: only groups of users authorized by the owner may view photos in the album


```

1  SPL Album {
2    Features = {Base, Remove, Owner, Groups, FineFit}
3    Deltas = {DBase, DRemove, DOwner, DRemoveAndOwner, DGroups,
4             DBaseFineFit, DRemoveFineFit, DOwnerFineFit, DGroupsFineFit}
5    Constraints { ... /* Same as in Listing 4 */ }
6    Partitions {
7      ... /* Same as in Listing 4 */
8      {DBaseFineFit} when (Base & FineFit);
9      {DRemoveFineFit} when (Remove & FineFit);
10     {DOwnerFineFit} when (Owner & FineFit);
11     {DGroupsFineFit} when (Groups & FineFit);
12   }
13   Products {
14     ... /* Same as in Listing 4 */
15     A.BaseFineFit = {Base, FineFit};
16     A.RemoveFineFit = {Base, Remove, FineFit};
17     A.OwnerFineFit = {Base, Owner, FineFit};
18     A.RemoveOwnerFineFit = {Base, Remove, Owner, FineFit};
19     A.OwnerGroupsFineFit = {Base, Owner, Groups, FineFit};
20     A.OwnerRemoveGroupsFineFit = {Base, Remove, Owner, Groups, FineFit};
21   }
22 }

```

Listing 11: Declaration of the Album PL extended with the products equipped with code for FineFit testing

As a result, the following operation table is generated:

login(n?:Name, p?:Password, report!:Report)	no loggedIn		some loggedIn
	n? in dom[users] and passwords[users[n?]] = p?	n? !in dom[users] or passwords[users[n?]] != p?	true
loggedIn	p? = ownerPassword	true	true
report!	loggedIn + users[n?]	=	=
	OK	AUTH_FAILED	ALREADY_IN

5.2 Delta-oriented programming of the retrieve function and of the driver

The retrieve function and the driver can be conveniently programmed in a delta-oriented way by extending the SPL with a new optional feature, FineFit, that when selected generates a product that is equipped with the retrieve method and the class extending the FineFitDriver base class (cf. Section 2.3).

To this aim, the declaration of the Album PL is modified by adding the feature FineFit and appending (after the last line of the partitions-block) the partitions with the delta modules for adding the retrieve function and the driver. The resulting declaration is given in Listing 11. There is no need for a FineFit delta module corresponding to RemoveAndOwner since it only changes the implementation of remove, and therefore does not affect the interface (cf. Listing 8).

The code of the delta module DBaseFineFit is given in Listings 12. The complete code of all the delta modules is available at [43].

¹⁶ Details can be found in Appendix B.

```

1  delta DBaseFineFit {
2
3      modifies it.unito.Album.ArrayPhotoAlbum {
4          import com.finefit.sut.IdMap;
5          import com.finefit.sut.State;
6          adds ... /* Same as in Listing 2 (bottom) */
7      }
8
9      adds { package it.unito.Album;
10         ... /* Same as in Listing 3 */
11     }
12
13 }

```

Listing 12: Delta module DBaseFineFit for adding the retrieve method and the FineFitDriver class to the Base product of the Album PL

5.3 Executing the tests

The product for a given feature configuration can be tested by: (i) generating its FineFit specification from the DeltaFineFit specification of the SPL, (ii) generating a version of the product equipped with the support code for using FineFit, and (iii) running FineFit. Note that the entire testing procedure for any product is fully automatic once the delta modules for implementing the feature FineFit have been written and the SPL declaration has been extended (cf. Section 5.2).

The number of products of an SPL can be exponential in the number of features. Testing all the products may be therefore unfeasible, even when the entire testing procedure is fully automatic (as in the case of DeltaFineFit). This issue can be handled by using *sample-based SPL testing* techniques [21,20,27,23], which address the problem of scalability by identifying a subset of products that is supposed to cover relevant combinations of features. Thus only the products in the identified subset have to be generated and tested by applying single-system testing techniques.

6 Evaluation of the approach

In this section we briefly evaluate the benefits of integrating data refinement based testing into delta-oriented SPL development from both a qualitative and a quantitative perspective.

6.1 Qualitative evaluation

The DeltaFineFit product-line declaration and the DeltaJ product-line declaration may be different in principle. For instance (as an extreme case), the complex-core DeltaFineFit specification of the Album PL can be used as a specification of the simple-core DeltaJ implementation of the Album PL (cf. Section 3.4 and 5.1). However, it is most effective to share the same product-line declaration for both the specification and the implementation, as illustrated in Section 5.1. Therefore, the delta-table construct must be expressive enough to allow delta-table modules to mimic delta

Features	Base		Base Remove		Base Owner		Base Remove Owner		Base Owner Groups		Base Remove Owner Groups	
	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>
Tables	1	1	1	1	1	1	1	1	1	1	1	1
Constants	1	1	1	1	1	1	1	1	1	1	1	1
Atoms	1	3	1	3	1	5	1	5	1	7	1	7
Enumeration	1	3	1	4	1	6	1	7	1	14	1	14
State	1	1	1	1	1	4	1	4	1	8	1	8
Invariants	1	2	1	2	1	3	1	3	1	6	1	6
Operations	3	8	4	11	5	18	6	2	10	53	11	57
TOTAL	8	18	9	22	10	37	11	42	15	89	16	93

Table 1 Size of the specifications of the 6 products of the Album PL without using DTMs. *T* is the number of tables and *L* is the total number of lines in these tables. The specification of the Base Album is illustrated in Fig. 3 and the specifications of the other 5 products are available at the DeltaFineFit home page [43].

modules, i.e., to allow developers to establish a one-to-one correspondence between delta-table modules and delta modules. Based on our experiments we believe that the delta-table construct presented in this paper satisfies this requirement.

Overall, the proposed integration of data refinement based testing with delta-oriented SPL development supports:

1. automatic generation of product specifications from delta-table modules that correspond to the delta modules by sharing the product-line declaration with the implementation;
2. automatic generation and inclusion of auxiliary code for FineFit’s refinement-based testing into the products (cf. Section 5.2); and
3. a clear layout of the commonalities and differences between product configurations in both the product specifications and in the auxiliary code needed for using FineFit.

6.2 Quantitative evaluation

In the following we present a quantitative evaluation of how concise is the specification of an SPL using the delta-oriented approach compared with the specification that does not use this approach. As case studies we considered the simple-core based and the complex-core based development of the Album PL. Table 1 illustrates the size of the specification of each of the six products of the Album PL, where each product is specified separately (without using the delta-oriented approach). In this counting scheme, we counted all rows of the tables excluding the title row for columns in constant, atom, state, and invariant tables.¹⁷ Using the information in Table 1 it is straightforward to compute for each kind of table the total number of tables and rows for all products (these numbers are presented in Table 2 in the columns of the naive approach).

Table 2 illustrates a quantitative comparison of three different specification approaches for the Album PL study:

¹⁷ The first row of enumeration tables, which consists of an atom name, was counted.

Approach	Naive		Simple-core delta-oriented			Complex-core delta-oriented		
Number of DTMs	6		5			4		
Tables	<i>T</i>	<i>L</i>	<i>T</i>	<i>L</i>	<i>in Figures</i>	<i>T</i>	<i>L</i>	The tables are available at the DeltaFineFit home page [43]
(Δ)Constants	6	6	1	1	3	1	1	
(Δ)Atoms	6	30	3	7	3, 7, 9	3	13	
(Δ)Enumeration	6	48	3	15	3, 7, 9	4	28	
(Δ)State	6	26	3	10	3, 7, 9	3	18	
(Δ)Invariants	6	22	3	6	3, 7, 9	3	12	
(Δ)Operation	39	169	11	75	3, 6, 7, 8, 9	26	103	
TOTAL	69	301	32	114	-	40	175	-

Table 2 Quantitative comparison of specification approaches. *T* is the number of tables and *L* is the total number of lines in these tables. For the naive specification approach, the specification of each of the 6 products of the Album PL can be considered as a core DTM (remember that a constants/atoms/enumeration/state/invariants table can be considered as a delta table).

- naive specification (i.e., each product is specified separately),
- simple-core delta-oriented specification, and
- complex-core delta-oriented specification.

The naive approach requires 69 tables with a total of 301 rows to specify the whole product line. Less than half the number of tables (32) and rows (114) are required when using DTMs in the simple-core delta-oriented specification approach. The number of tables and rows in the complex-core delta-oriented specification approach is bigger than in the simple-core approach. Note that:

- DTMs in simple-core specifications define how to add and expand tables,
- DTMs in complex-core specifications define how to remove and trim tables, and
- the set of delta tables for expanding tables in the simple-core specification has roughly the size (in the number of both tables and rows) of the set of delta tables for removing and trimming tables in the complex-core specification.

Therefore the difference in the number of tables and rows is mainly due to the fact that the DTM used to build the core for complex-core specifications is larger than the one used for simple-core specifications, in the number of both tables (16 vs 8) and rows (93 vs 18)—c.f. Table 1. Conversely, the complex-core specification has less DTMs.

It is worth to observe that the ability to support complex-core specifications and, more general, the ability to build specifications where the DTMs mimic the delta modules in the implementation (cf. Section 6.1) makes it possible to use DeltaFineFit in reactive and extractive SPL development (cf. Section 3.4).

The case study confirms our expectations that the delta-oriented specification approach reduces the size of the specification of the SPL (both with the simple- and complex-core approaches). The Album PL has only 4 features and 6 products. We expect that these benefits will be even more significant in more complex software product lines.

7 Related Work

The use of a tabular notation to specify operations was inspired by Parnas [32]. Tabular expressions can be represented and interpreted in many ways. Parnas proposes precise definitions for ten varieties of tabular expressions [32]. We use the simple and intuitive Vector Function Tables to describe functions.

Various surveys deal with testing in SPLs; e.g., [16, 30, 25, 8]. Delta tables leverage the idea of tabular notation to the specification of product lines. To the best of our knowledge this idea has not been proposed yet. The only papers introducing similar operations on tables, also named delta tables, describe updates or views on database, e.g. [35]. However, it seems that this idea has not been proposed before to specify programs.

The approach proposed by Uzuncaova et al. [42] seems the most related to our testing approach. The paper describes a scope-bounded testing technique for programs in an SPL developed by using the feature oriented programming (FOP) approach (cf. Section 1 and Section 3.4). In the FOP approach a program is a composition of feature modules, where each module implements an increment in the functionality of the program corresponding to a feature. This testing approach differs from ours as their test suite enables only to add features, whereas we allow also to remove features. Hence, the test suite in [42] is monotonically extended with each additional feature. Using the Alloy Analyzer, tests are generated and the correctness of the program is checked. Each feature F is specified as an Alloy formula, and the Alloy specification of a product is the conjunction of the specifications for its features. To generate larger inputs, tests are generated incrementally, as the analyzer is used each time on partial specifications rather than on the complete specification of the program. First the analyzer solves the specification for the Base feature (which corresponds to the base product), and the created relations are used when running the analyzer on the specification of a feature F that is composed with Base. In our approach, we generate both the program and its specification using the DOP approach, where a product is generated from another one by applying a delta-operation that may add or remove functionalities. The tests are then generated by considering the complete specification of a product. However, when a program is a composition of features, i.e., when the delta-modules do not use the remove operation, it should be possible to apply our approach by using the incremental approach as in [42].

More recently, Lochau et al. [29, 28] proposed a model-based SPL testing framework that is based on a delta-oriented SPL test model and on regression-based test artifact generation. The framework comprises state machines as test models extended with delta-oriented modeling concepts to express variability. The framework is aimed to capture reuse potential, reusing both (i) test cases for different product variants, and (ii) test results across different products. In our framework no reuse of test cases or test results is currently supported. We believe that the technique proposed in [28] could be applicable also to our approach, although we are using a different specification language (Parnas tables).

Proof systems for the verification of delta-oriented SPLs have been recently proposed [18, 11]. These approaches are related to ours since they use DOP to generate a product together with its specification; however, they aim at a formal verification.

Hähnle and Schaefer [18] presented a verification approach for DOP that relies on the Liskov principle for DOP, i.e., the specification of a method modified by a delta module must entail the specification of the previous versions of the method. This ensures that each delta module can be verified by approximating called methods that are defined in other delta modules by the specification of their first introduction. In order to relax the restriction imposed by the Liskov principle for DOP [18], Damiani et-al. [11] suggest to use symbolic assumptions on called methods in order to separate the specifications of method implementations from the requirements to method calls in a way which is similar to lazy behavioral subtyping [15]. The idea is to verify first each delta module in isolation, based on symbolic assumption for methods calls that may be defined in other delta modules, and then to verify each product based on the already established specification for the used delta modules. The combination of software verification and testing techniques is encouraged due to their complementary strengths. We believe that some ideas for combining verification and testing that already had been implemented for a single system (see [6]) could be applied to SPLs.

8 Conclusion and Future Work

We illustrated how the FineFit approach can be integrated into DOP to support the development of correct software product lines. Refinement-based testing is completely integrated into the delta-oriented SPL development process. Delta-oriented programming is used not only to generate the system under test, but also to generate the FineFit specification, the retrieve function and the driver for each product.

We would like to complete the development of the DeltaFineFit tool chain. In particular, we would like to improve the tool support for the delta-oriented programming of the retrieve function and the driver by developing a tabular notation that will be automatically translated into the corresponding delta modules.

A type system for DOP, formalized for the minimal core calculus IMPERATIVE FEATHERWEIGHT DELTA JAVA(IFΔJ) [7], guarantees that if an IFΔJ product line is well-typed then its product generation mapping is total and all its products are well-typed Java programs (cf. Section 3.3). In future work we plan to formalize a suitable notion of well-formed FineFit specification and to develop static analysis techniques for checking whether the product generation mapping of a DeltaFineFit specification is total and all its products are well-formed FineFit specifications. Moreover, we would like to extend the DeltaFineFit approach to deal with dynamic DOP [13, 12].

We also plan to develop case studies to assess the effectiveness of using DeltaFineFit with sample-based SPL testing techniques (see Section 5.3), and to enhance DeltaFineFit by exploring both the use of techniques for reusing test cases or test results similar to [28] (cf. Section 7) and the use of test prioritization techniques similar to [26]. Another enhancement of DeltaFineFit that we would like to explore is the integration of verification techniques (see the discussion at the end of Section 7).

Acknowledgements We thank the anonymous referees of PPPJ'13 for valuable comments on a preliminary version of this paper and the anonymous SoSyM referees for many insightful comments and suggestions for improving the paper.

A The structure and semantics of operation tables

A FineFit operation table is a predicate that specifies the behavior of an operation as a relation between the model's state variables before the operation starts (the pre state) and after the operation completes (the post state). It consists of two major areas: an expression table and a precondition tree. The expression table is a set of columns, where each column defines the values of the state variables in the post state, given their values in the pre state. The precondition tree consists of predicates that determine which columns to use in the definition of the post state. For example, consider the following operation table:

R($c?:\text{Int}$)	$c? \leq 0$		$c? > 0$
	$x < 0$	$x \geq 0$	true
x	=	0	$x + c?$
y	=	0	$x + y$

The operation R has a single input parameter $c?$ and determines the value of its two state variables x and y as follows: If $c?$ is positive then the operation must set x to $x + c?$ and y to $x + y$ (the right most column), otherwise, if $c?$ is not positive, the effect of the operation depends on the current value of x . When x is not negative x and y must be set to zero and when x is negative then x and y do not change (their value must be the same as it was when the operation started). Note that the predicate for determining whether columns 1 and 2¹⁸ are relevant for the specification is the conjunction of the child predicates $x < 0$ and $x \geq 0$ with their parent predicate $c? \leq 0$.

We now define the conditions necessary for the precondition part of the table to form a tree. Consider the precondition part of the table as a matrix of cells. Then, each predicate spans one or more consecutive cells (in the same row) and the space that each predicate occupies must be included below the space that its parent (the predicate above) occupies. For example, the predicate $x \geq 0$ occupies cell (2, 2) and its parent, the predicate $c? \leq 0$ occupies cells (1, 1), (1, 2).

Formally, consider a precondition part that is arranged in a matrix of n rows by m columns of cells. The i -th row contains a sequence of predicates, each occupying a span of one or more cells. Let k_i indicate the sequence of spans that each predicate occupies (and therefore k_{ij} is the j -th span in the i -th row). For example, in the operation above, $k_1 = 2, 1$ and $k_2 = 1, 1, 1$. In order for the spans to describe a valid precondition tree, three conditions must be met:

1. The sum of spans in each row must be equal to the number of columns in the matrix:

$$\sum_{j=1}^{|k_i|} k_{ij} = m$$

where $|k_i|$ is the number of elements in the i th sequence of spans.

2. The cells that each predicate spans must be included below the span of cells of its parent:

$$S_{ij} \subseteq S_{(i-1)l} \quad \text{for some } 1 \leq l \leq m$$

where the span of cells of the j -th predicate in the i -th row is

$$S_{ij} = \sum_{l=1}^{j-1} k_{il} + 1, \dots, \sum_{l=1}^j k_{il}$$


3. The last row (n) must consist of individual cells, that is

$$k_{nj} = 1 \quad \text{for all } 1 \leq j \leq m.$$

This arrangement ensures that the predicates form a tree structure¹⁹ with m leaves, all occupying the last row. The following matrix illustrates this structure:

¹⁸ We do not count the left most column containing the state variables.

¹⁹ In fact this is a forest. However, it can be transformed into a tree by adding an imaginary *true* as the parent of the predicates at the first row.



There are three span sequences, one for each row:

$$k_3 = 1, 1, 1, 1, 1, 1, 1 \quad (3)$$

$$k_2 = 3, 2, 2 \quad (4)$$

$$k_1 = 5, 2 \quad (5)$$

We can see that each span is included below the span of its parent. For example,

$$S_{2,2} = \{4, 5\} \subseteq S_{1,1} = \{1, 2, 3, 4, 5\}$$

We say that the precondition tree is well formed when it satisfies the three conditions we have defined above. In the rest of the discussion we always assume that we are working with well formed precondition trees. Given a column i we define the guard for this column as the conjunction of the leaf that occupies the i -th column and all its ancestors. Let $\vec{v}' = v'_1, \dots, v'_l$ be the state space vector. Let \vec{c}_i be the i -th column of the expression table. Then the meaning of the operation specification is:

$$\vec{v}' \in \{\vec{c}_i : 1 \leq i \leq m \wedge \text{guard}(i)\}$$

That is, the value of the state variables vector in the post state can be equal to the value of any of the expression table columns whose guard was true in the pre state.

The semantics of individual predicates and expressions is explained in [17].

B The algorithm that computes the *apply* function

We provide (in Section B.1) a Haskell executable specification of the *apply* function (cf. Section 4.2.1), and present (in Section B.2) its Java implementation. Note that in this paper all the tables resulting from delta-table application have been generated using this implementation.

B.1 Executable specification in Haskell

The executable specification of the algorithm that computes the *apply* function is provided as a program written in Haskell (Listing 13). Haskell is a functional programming language with a syntax that strongly resembles the usual mathematical notation for defining function by cases, via pattern matching.

The Haskell code and the comments in Listing 13 should be almost self-explanatory. In the following, we shortly explain some technicalities. The expression contained in an ordinary table cell and in a delta table cell containing the insert or replace operators is represented using the standard library type **String** (Line 4). Both ordinary tables and delta tables are represented as trees, where each node corresponds to a cell. A tree that represents a table (ordinary or delta) is expressed as a value of the recursive data type **Table** (Line 11)—a value of type **Table** describes the root cell and the list of its immediate subtrees. The data type **Table** has five data constructors (each describes a different kind of node): **Basic**, **Insert**, and **Replace**

(all have arity two), **Match** and **Remove** (arity one).²⁰ Ordinary tables are represented by using only the data constructor **Basic**, while delta tables are represented by using only the other data constructors, which correspond to the delta table operators.

Example 3 (Representation of ordinary tables and delta tables in Haskell) Since ordinary tables may have more than one cell in their first row, they are always encoded by adding a top node containing the string "ROOT". Delta tables are therefore encoded by adding a top node containing the match operator (matching the top node in the ordinary tables). The following ordinary and delta tables

A	
C	D
1	2
i	ii

*			
►I		►J	
—	*	*	—
—	*	*►X	—
—	*	*►Y	—

are respectively represented by the following values of type **Table**:

```
Basic "ROOT" [Basic "A" [Basic "C" [Basic "1" [Basic "i" []],
                               Basic "D" [Basic "2" [Basic "ii" []]]],
              Match [Match [Insert "I" [Remove [Remove [Remove []],
                                                    Match [Match [Match []]],
                                                    Insert "J" [Match [Replace "X" [Replace "Y" []],
                                                                    Remove [Remove [Remove []]]]]],
                      Remove [Remove [Remove []]]]]]
```

The Haskell function **apply** (Lines 16-51) corresponds to the *apply* function (introduced in Section 4.2.1) that executes the delta operations of the delta table. To improve readability we have not modeled the behavior described in Remark 2 of Section 4.2.1.

Line 26 declares the type of the **apply** function. The first argument of the function **apply** has type **Maybe Table**. The standard library **Maybe** data type has two constructors: the unary data constructor **Just** and the constant **Nothing**. It is used to specify optional values: a value of type **Maybe Table** either contains a table **t** (represented as **Just t**), or it is empty (represented as **Nothing**).

The standard library function **concat** takes two lists and returns their concatenation. The standard library function **zipWith** calls a given function pairwise on each member of both lists, returning a list. For the convenience of readers **zipWith**'s code is as follows:

```
zipWith :: (a -> b -> c) -> [a] -> [b] -> [c]
zipWith _ [] = []
zipWith _ [] = []
zipWith f (x:xs) (y:ys) = (f x y) : zipWith f xs ys
```

*Example 4 (An execution of the Haskell function **apply**)* Consider the following delta table application, which already has been presented at the end of Section 4.2.1 (recall that the numbers to the left of the cells denote the recursion step):

1:	A		1:	*		1:	A
2,3:	C		2:	►I		2:	I
4:	1		3:	*		3:	C
5:	i		4:	*►D		4:	D
6:	w		5:	—		5:	w
			6:	*		6:	

The above delta table application can be defined by appending to the code in Listing 13 the following three lines:²¹

```
t = Basic "ROOT" [Basic "A" [Basic "C" [Basic "1" [Basic "i" [Basic "w" []]]],
                  Match [Match [Insert "I" [Match [Remove [Insert "D" [Remove [Match []]]]]],
                          Remove [Remove [Remove []]]]]]
main = print (apply Nothing t dt)
```

Executing the Main program yields

```
Basic "ROOT" [Basic "A" [Basic "I" [Basic "C" [Basic "D" [Basic "w" []]]]]]
```

²⁰ Adding **deriving(Show)** at the end of a data declaration enables printing instances on standard output.

²¹ Adding a definition for the name **main** makes Listing 13 a complete program.

```

1  module Main where
2
3
4  type Expr = String
5
6
7  -- |A table is a tree that consists of the following nodes:
8  -- 1. Basic nodes that represent the expressions in Finefit tables.
9  -- 2. Match, Remove, Insert, and Replace nodes, that represent the operations of the delta tables.
10
11 data Table = Basic Expr [Table]
12            | Match [Table] | Remove [Table] | Insert Expr [Table] | Replace Expr [Table]
13            deriving(Show)
14
15
16 -- |We assume that the tables given to the apply function, have a correct structure. For example, that each
17 -- table is a tree rooted at a dummy node, the original table has only basic expression nodes, and so on.
18 -- The complete system checks that the tables have the correct structure, adding these checks to the
19 -- specification would hinder understanding the essential ideas.
20 -- The apply function takes the following arguments:
21 -- 1. the parent of the node in the original table (or the constant Nothing, if the node is the root),
22 -- 2. the node itself,
23 -- 3. the corresponding node in the delta table.
24 -- The function returns the sequence of nodes that corresponds to the new table.
25
26 apply :: Maybe Table -> Table -> Table -> [Table]
27
28 -- |Some notes on the recursive specification of the function:
29 -- 1. We divide the specification into several cases by pattern matching on the kinds of delta table nodes
30 --    (each kind of delta table operator node has its own case).
31 -- 2. In all cases, we take care to pass the correct parent for the nodes in the recursive call.
32 -- Applying a delta table match operator to a node of the original table, keeps the content of the original
33 -- table, and recursively applies the children of the match node, to those of the original node.
34
35 apply _ node (Match deltas) = [Basic (content node) (apply_to_children node deltas)]
36
37 -- |Applying a delta table remove to a node of the original table, replaces the node by its children
38 -- (after they are transformed by the children of the delta table node).
39 -- This case is the reason why apply returns a list of tables.
40
41 apply _ node (Remove deltas) = apply_to_children node deltas
42
43 -- |Applying a delta table insert to a node of the original table, inserts a new node as the child of the
44 -- original node's parent, and all of the parent's children (including the original node) as the children of
45 -- the newly inserted node.
46
47 apply (Just parent) _ (Insert e' deltas) = [Basic e' (apply_to_children parent deltas)]
48
49 -- Replace is implemented in terms of remove followed by insert
50
51 apply p x (Replace e cs) = apply p x (Insert e [Remove cs])
52
53
54 -- |To process the children of a table node against a sequence of delta table operators, we traverse both
55 -- lists, calling apply on each pair of corresponding nodes. This generates a list of lists, that we then
56 -- concatenate to form the result.
57
58 apply_to_children parent deltas = concat (zipWith (apply (Just parent)) (children parent) deltas)
59
60 children :: Table -> [Table]
61 children (Basic _ xs) = xs
62
63 content :: Table -> Expr
64 content (Basic c _) = c

```

Listing 13: Executable specification of the algorithm that computes the *apply* function introduced in Section 4.2.1; the behavior described in Remark 2 is not modeled

B.2 The Java implementation

The Java implementation of the algorithm that computes the *apply* function is given in Listings 14 and 15. A cell of a table is represented by the *Node* class which extends the *Vector* class. Tables are represented as trees, and an object of class *Node* represents a cell, which is the root of a subtree, and contains (in the *Vector*'s elements) the reference to the roots of the immediate subtrees.


The *applyPrime* method (Lines 16–21 in Listing 14) corresponds to the *apply'* function (introduced in Section 4.3). It first calls (Line 17) the *prepare* method (which applies the rules described in Section 4.3), and then invokes (Line 19) the *apply* method (Lines 29–58 in Listing 14). The *apply* method corresponds to the *apply* function (introduced in Section 4.2.1 and specified in Section B.1). When calling the *apply* method, the *this* object is the current node of the delta table, and *orig* and *res* are the current nodes of the original table and of the resulting table, respectively.

The body of the *apply* method implements the recursive walking of the parse tree, which is controlled by the operations in the cells of the delta table and, therefore, describes how each operation works. The first part of the method (Lines 31–44) reads the operation of the current delta node and updates the content of the resulting node; Lines 35–40 implement the behavior described in Remark 2 of Section 4.2.1. The second part of the method (Lines 46–56) is responsible for the recursive call.²² At each recursive call (Lines 37, 50, and 54) the method *getOrCreateChild(i)*, which creates the nodes of the resulting table, is called; it either returns the *i*-th subnode if it already exists, or it creates the *i*-th subnode and returns it. We now continue illustrating the *apply* method using small examples. (An example of an execution trace of the method *apply* is provided in Example 5.)

The loop that begins in Line 46 iterates over the siblings of the current delta table node. This means that if at the current level the original table has more cells than the delta table, then the additional cells are ignored. For example:




If the delta table has more cells than the original table at the current level, the last cell of the original table is repeatedly used when processing the additional operations of the delta table (Line 53). For example:



If the original table contains no cells at the current level and the current operation of the delta table is *** or *—*, the remaining operations of the delta table on the current branch are ignored (Line 52).



If the original table contains no cells at the current level and the current operation of the delta table is *►* or **►*, a cell with the value of the insert or replace operations is inserted (Line 50) in the resulting table. Line 47 maps in this case the replace operation to an insert operation. For example:



It is therefore possible to extend a table with new variable and value rows by using either the insert or replace operation.

The algorithm traverses the original table and the delta table in parallel. The recursive call in Line 54 is responsible for the operations match, remove, and replace²³. The recursive call in Line 50 handles the *insert* operation, where the current node of the original table is passed as the first argument rather than the current child of the original table. Delaying the recursive step on the original table results in the described semantics of the insert operation.

The recursive call in Line 37 is used for convenience. If either the remove or the match operation occur as the last condition in a condition hierarchy, but the condition cell *c* of the original table has additional subconditions (this is checked by *isLastConditionNode()*), the remove or the match operation is applied to all subcondition of *c*. For example, let *A*, *B*, and *C* be condition cells, then

²² Line 37 is a special case, where *** matches multiple condition cells.

²³ If *replace* is not substituted by *insert* in Line 47.

```

1  import java.util.Vector;
2
3  /* The Node class represents a cell of a table. Subnodes represent cells below this cell. */
4
5  public class Node extends Vector<Node>{
6
7      public enum Op {nop,match,remove,insert,replace};
8      public Op op = Op.nop; //Default is no operation
9      public String val = ""; //Value or content of the cell
10     public boolean cond = false; //true if this node is a condition node and not a value node
11
12     public Node(String val){this.val=val;}
13
14
15     /* "this" represents the root of the delta table. "orig" is the root of the original table. */
16     public Node applyPrime(Node orig){ //corresponds to the function apply' of Section 4.3
17         prepare(orig, this);
18         Node result = new Node("<root>");
19         apply(orig, result);
20         return result;
21     }
22
23
24     public static void prepare(Node orig, Node delta){...} //See Section 4.3
25
26
27     /* The parameters "this", "orig", and "res" represent the current node (or cell) of the
28        delta table, original table, and resulting table, respectively. */
29     public void apply(Node orig, Node res){ //corresponds to the function apply of Section 4.2.1
30
31         switch(op){
32             case nop: break;
33             case remove: res.op = Op.remove; //no break here
34             case match: res.val = unifyValue(orig.val);
35                 if(isLastConditionNode() && !orig.isLastConditionNode()){ //see Remark 2 in Section 4.2.1
36                     for(int i=0; i<orig.size(); i++){
37                         apply(orig.get(i), res.getOrCreateChild(i, cond));
38                     }
39                     return;
40                 }
41             case insert: res.val = val; break;
42             case replace: res.val = val; break;
43         }
44
45         for(int i=0; i<size(); i++){
46             if(orig.size()==0 && get(i).op==Op.replace){ get(i).op=Op.insert; }
47
48             if(get(i).op==Op.insert){
49                 get(i).apply(orig, res.getOrCreateChild(i, cond));
50             }else{
51                 if(orig.size()==0) break;
52                 Node newOrig = (i<orig.size())? orig.get(i) : orig.lastElement() );
53                 get(i).apply(newOrig, res.getOrCreateChild(i, cond));
54             }
55         }
56         res.deleteTemporaryRemoveNodes();
57     } // end of method apply
58
59
60     ... //See Listing 15
61
62 } // end of class Node
63

```

Listing 14: Java implementation of the algorithms that compute the *apply'* and *apply* functions

```


1  /* Returns true if this node is a condition node and its child nodes are non-condition (value) nodes.*/
2  public boolean isLastConditionNode(){
3      if(!cond) return false;
4      if(size()==0) return true;
5      if(get(0).cond) return false; //For simplicity only one child is checked.
6      return true;
7  }
8
9  /* Returns the i'th subnode. If the subnode does not exist, then it is create and returned. */
10 public Node getOrCreateChild(int i, boolean cond){
11     if(i<size()){
12         return get(i);
13     }else{
14         Node tmp = new Node("");
15         tmp.cond = cond;
16         tmp.op = Op.nop;
17         add(tmp);
18         return getOrCreateChild(i, cond);
19     }
20 }
21
22 /* Replaces each child of this node that is marked as remove "-" by its grand children.
23    The algorithm operates recursively until no child is marked with remove "-".*/
24 protected void deleteTemporaryRemoveNodes(){
25     boolean childRemoved=false;
26     do{
27         childRemoved=false;
28         for(int i=0; i<size();i++){
29             Node child = get(i);
30             if(child.op==Node.Op.remove){ //Replace the removed child by its children
31                 remove(i);
32                 for(int j=i,k=0; k<child.size();k++,j++){
33                     insertElementAt(child.get(k), j);
34                 }
35                 childRemoved = true;
36                 break;
37             }
38         }
39     }while(childRemoved);
40 }
41
42 /* if the delta value is "*", then origVal is returned. Otherwise all occurrences
43    of "*" in the delta value are replace with origVal. E.g., "x" apply "(*)+y" yields "x+y".*/
44 public String unifyValue(String origVal){
45     if(val.equals("*"))
46         return origVal;
47     else
48         return val.replaceAll("(*)",origVal);
49 }

```


Listing 15: Helper methods of the apply method



The reader may have noticed that unlike the other operations, the remove operation is not handled in the loop in Lines 46–56. When a remove operation occurs, a temporary cell is created in the resulting table that is marked to be removed (Line 33). The marked cells are removed in Line 57. Removing nodes at the end of the recursion simplifies the implementation, because removing a node c pulls-up its children c_1, \dots, c_n to the current level. Hence, by removing a node the number of siblings at the current level may decrease, stay constant, or increase. In the following example we assume that A is not a condition cell.



Example 5 (An execution trace of the method apply) For a complete example consider the following delta table application, which already has been presented at the end of Section 4.2.1 and in Example 4 of Appendix B.1:



Executing the method `apply` in (Lines 29–58) will result in the following recursive invocations of the method:

Recursion step	this	orig	res
0	< root >	< root >	< root >
1	*	A	A
2	►I	C	I
3	*	C	C
4	*►D	I	D
5	—	i	—
6	*	w	w

At the end of the last recursion step of the `apply` algorithm (Line 57 in Listing 14) the method `deleteTemporaryRemoveNodes` is called which removes the temporary “—” node from the resulting tree `res` that is introduced at recursion step 5.

References

1. Apel, S., Batory, D.S., Kästner, C., Saake, G.: Feature-Oriented Software Product Lines: Concepts and Implementation. Springer (2013)
2. Apel, S., Kästner, C., Grösslinger, A., Lengauer, C.: Type safety for feature-oriented product lines. *Automated Software Engineering* **17**(3), 251–300 (2010)
3. Apel, S., Kästner, C., Lengauer, C.: Feature featherweight java: A calculus for feature-oriented programming and stepwise refinement. In: *Proceedings of the 7th International Conference on Generative Programming and Component Engineering, GPCE ’08*, pp. 101–112. ACM, New York, NY, USA (2008). DOI 10.1145/1449913.1449931
4. Batory, D.: Feature Models, Grammars, and Propositional Formulas. In: *Proceedings of the International Conference on Software Product Lines (SPLC), LNCS*, vol. 3714, pp. 7–20. Springer (2005)
5. Batory, D., Sarvela, J., Rauschmayer, A.: Scaling Step-Wise Refinement. *IEEE TSE* **30**(6), 355–371 (2004)
6. Beckert, B., Gladisch, C., Tyszberowicz, S., Yehudai, A.: KeYGenU: combining verification-based and capture and replay techniques for regression unit testing. *Int. J. Systems Assurance Engineering and Management* **2**(2), 97–113 (2011)
7. Bettini, L., Damiani, F., Schaefer, I.: Compositional type checking of delta-oriented software product lines. *Acta Informatica* **50**, 77–122 (2013)
8. do Carmo Machado, I., McGregor, J.D., Cavalcanti, Y.C., de Almeida, E.S.: On strategies for testing software product lines: A systematic literature review. *Information and Software Technology* **56**(10), 1183 – 1199 (2014)
9. Clements, P., Northrop, L.: *Software Product Lines: Practices and Patterns*. Addison Wesley Longman (2001)
10. Damiani, F., Gladisch, C., Tyszberowicz, S.: Refinement-based testing of delta-oriented product lines. In: *Proceedings of the 2013 International Conference on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools, PPPJ ’13*, pp. 135–140. ACM, New York, NY, USA (2013). DOI 10.1145/2500828.2500841

11. Damiani, F., Owe, O., Dovland, J., Schaefer, I., Johnsen, E.B., Yu, I.C.: A transformational proof system for delta-oriented programming. In: Proceedings of the 16th International Software Product Line Conference - Volume 2, SPLC '12, pp. 53–60. ACM, New York, NY, USA (2012). DOI 10.1145/2364412.2364422
12. Damiani, F., Padovani, L., Schaefer, I.: A formal foundation for dynamic delta-oriented software product lines. In: Proceedings of the 11th International Conference on Generative Programming and Component Engineering, GPCE '12, pp. 1–10. ACM, New York, NY, USA (2012). DOI 10.1145/2371401.2371403
13. Damiani, F., Schaefer, I.: Dynamic delta-oriented programming. In: Proceedings of the 15th International Software Product Line Conference, Volume 2, SPLC '11, pp. 34:1–34:8. ACM, New York, NY, USA (2011). DOI 10.1145/2019136.2019175
14. Delaware, B., Cook, W.R., Batory, D.: Fitting the pieces together: A machine-checked model of safe composition. In: Proceedings of the the 7th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on The Foundations of Software Engineering, ES-EC/FSE '09, pp. 243–252. ACM, New York, NY, USA (2009). DOI 10.1145/1595696.1595733
15. Dovland, J., Johnsen, E.B., Owe, O., Steffen, M.: Lazy behavioral subtyping. *Journal of Logic and Algebraic Programming* **79**(7), 578–607 (2010)
16. Engström, E., Runeson, P.: Software product line testing - a systematic mapping study. *Information and Software Technology* **53**(1), 2–13 (2011)
17. Faitelson, D., Tyszbrowicz, S.S.: Data refinement based testing. *Int. J. Systems Assurance Engineering and Management* **2**(2), 144–154 (2011)
18. Hähnle, R., Schaefer, I.: A Liskov Principle for Delta-Oriented Programming. In: Leveraging Applications of Formal Methods, Verification and Validation. Technologies for Mastering Change International Symposium (ISoLA), Part I, *LNCS*, vol. 7609, pp. 32–46. Springer (2012)
19. Jackson, D.: *Software Abstractions - Logic, Language, and Analysis*. MIT Press (2012)
20. Johansen, M.F., Haugen, O., Fleurey, F.: Properties of realistic feature models make combinatorial testing of product lines feasible. In: Proceedings of the International Conference on Model Driven Engineering Languages and Systems (MODELS), pp. 638–652. Springer-Verlag, Berlin, Heidelberg (2011)
21. Johansen, M.F., Haugen, O., Fleurey, F.: An algorithm for generating t-wise covering arrays from large feature models. In: Proceedings of the 16th International Software Product Line Conference - Volume 1, SPLC '12, pp. 46–55. ACM, New York, NY, USA (2012). DOI 10.1145/2362536.2362547
22. Koscielny, J., Holthusen, S., Schaefer, I., Schulze, S., Bettini, L., Damiani, F.: Deltaj 1.5: Delta-oriented programming for java 1.5. In: Proceedings of the 2014 International Conference on Principles and Practices of Programming on the Java Platform: Virtual Machines, Languages, and Tools, PPPJ '14, pp. 63–74. ACM, New York, NY, USA (2014). DOI 10.1145/2647508.2647512
23. Kowal, M., Schulze, S., Schaefer, I.: Towards efficient spl testing by variant reduction. In: Proceedings of the 4th International Workshop on Variability & Composition, VariComp '13, pp. 1–6. ACM, New York, NY, USA (2013). DOI 10.1145/2451617.2451619
24. Krueger, C.: Eliminating the Adoption Barrier. *IEEE Software* **19**(4), 29–31 (2002)
25. Lee, J., Kang, S., Lee, D.: A survey on software product line testing. In: Proceedings of the 16th International Software Product Line Conference - Volume 1, SPLC '12, pp. 31–40. ACM, New York, NY, USA (2012). DOI 10.1145/2362536.2362545
26. Li, Z., Harman, M., Hierons, R.M.: Search algorithms for regression test case prioritization. *IEEE TSE* **33**(4), 225–237 (2007)
27. Lochau, M., Goltz, U.: Feature interaction aware test case generation for embedded control systems. *Electronic Notes in Theoretical Computer Science* **264**(3), 37–52 (2010)
28. Lochau, M., Lity, S., Lachmann, R., Schaefer, I., Goltz, U.: Delta-oriented model-based integration testing of large-scale systems. *Journal of Systems and Software* **91**(0), 63 – 84 (2014)
29. Lochau, M., Schaefer, I., Kamischke, J., Lity, S.: Incremental model-based testing of delta-oriented software product lines. In: TAP, *LNCS*, vol. 7305, pp. 67–82. Springer (2012)
30. da Mota Silveira Neto, P.A., do Carmo Machado, I., McGregor, J.D., de Almeida, E.S., de Lemos Meira, S.R.: A systematic mapping study of software product lines testing. *Information and Software Technology* **53**(5), 407 – 423 (2011). Special Section on Best Papers from {XP2010}
31. Mugridge, R., Cunningham, W.: *Fit for Developing Software: framework for integrated tests*. Prentice Education, Inc. (2005)
32. Parnas, D.L.: Tabular representation of relations. Tech. Rep. 260, Research Institute of Ontario, McMaster University (1992)

33. Pohl, K., Böckle, G., van der Linden, F.: Software Product Line Engineering - Foundations, Principles, and Techniques. Springer (2005)
34. de Roeper, W.P., Engelhardt, K.: Data Refinement: Model-oriented Proof Theories and their Comparison, *Cambridge Tracts in Theoretical Computer Science*, vol. 46. Cambridge University Press (1998)
35. Salem, K., Beyer, K., Lindsay, B., Cochrane, R.: How to roll a join: Asynchronous incremental view maintenance. In: Proceedings of the 2000 ACM SIGMOD International Conference on Management of Data, SIGMOD '00, pp. 129–140. ACM, New York, NY, USA (2000). DOI 10.1145/342009.335393
36. Schaefer, I., Bettini, L., Bono, V., Damiani, F., Tanzarella, N.: Delta-oriented Programming of Software Product Lines. In: Software Product Line Conference (SPLC), *LNCIS*, vol. 6287, pp. 77–91. Springer (2010)
37. Schaefer, I., Bettini, L., Damiani, F.: Compositional type-checking for delta-oriented programming. In: Proceedings of the Tenth International Conference on Aspect-oriented Software Development, AOSD '11, pp. 43–56. ACM, New York, NY, USA (2011). DOI 10.1145/1960275.1960283
38. Schaefer, I., Damiani, F.: Pure delta-oriented programming. In: Proceedings of the 2Nd International Workshop on Feature-Oriented Software Development, FOSD '10, pp. 49–56. ACM, New York, NY, USA (2010). DOI 10.1145/1868688.1868696
39. Schaefer, I., Rabiser, R., Clarke, D., Bettini, L., Benavides, D., Botterweck, G., Pathak, A., Trujillo, S., Villela, K.: Software diversity: state of the art and perspectives. *International Journal on Software Tools for Technology Transfer* **14**(5), 477–495 (2012)
40. Spivey, J.M.: The Z Notation: A Reference Manual. Prentice Hall International (2001)
41. Torlak, E., Jackson, D.: Kodkod: A relational model finder. In: Tools and Algorithms for the Construction and Analysis of Systems (TACAS), *LNCIS*, vol. 4424, pp. 632–647. Springer (2007)
42. Uzuncaova, E., Khurshid, S., Batory, D.S.: Incremental test generation for software product lines. *IEEE TSE* **36**(3), 309–322 (2010)
43. The DeltaFineFit home page. <http://di.unito.it/deltafinefit>
44. The DeltaJ home page. <https://www.tu-braunschweig.de/isf/research/deltas>
45. The FineFit home page. <https://github.com/coderocket/finefit>