

Document downloaded from:

<http://hdl.handle.net/10251/71738>

This paper must be cited as:

Silva, J.; Insa Cabrera, D. (2015). A Generalization and Paradigm-Independent Reformulation of Algorithmic Debugging. En Logic-Based Program Synthesis and Transformation. Springer. 261-276. doi:10.1007/978-3-319-27436-2_16.



The final publication is available at

http://link.springer.com/chapter/10.1007/978-3-319-27436-2_16

Copyright Springer

Additional Information

The final publication is available at Springer via http://dx.doi.org/10.1007/978-3-319-27436-2_16

A Generalization and Paradigm-Independent Reformulation of Algorithmic Debugging

David Insa and Josep Silva

Universitat Politècnica de València, Valencia, Spain
{dinsa, jsilva}@dsic.upv.es

Abstract. Algorithmic debugging is a semi-automatic debugging technique that is present in practically all mature programming languages. In this paper we claim that the state of the practice in algorithmic debugging is a step forward compared to the state of the theory. In particular, we argue that novel techniques for algorithmic debugging cannot be supported by the standard internal data structures used in this technique, and a generalization of the standard definitions and algorithms is needed. We identify two specific problems of the standard formulation and implementations of algorithmic debugging, and we propose a reformulation to solve both problems. The reformulation has been done in a paradigm-independent manner to make it useful and reusable in different programming languages.

1 Introduction

One of the most important semi-automatic debugging techniques is algorithmic debugging (AD) [27]. This technique has experienced a significant advance in the last decade. Concretely, new techniques have been proposed to improve performance [9, 15], to improve scalability [11], to improve interaction with the user [6], and to improve GUIs [12, 13]. The maturity of these techniques has eventually lead to the integration of algorithmic debuggers into sophisticated programming environments. Two interesting cases are [12] and [11], which combine AD with the standard debugging perspective of Eclipse [1]. The main advantage of AD is its high level of abstraction. It is even possible to debug a program without looking at the source code.

Example 1. Let us assume the existence of a buggy Java code composed of three methods: `int add(int x, int y)` sums its two arguments, `boolean isEven(int x)` returns `true` if its only argument is even, or `false` otherwise; and, `int sumNumbers(int[] array, String eo)` takes an array of integers and sums the elements that are even or odd depending on the value of the second argument. Therefore, with the following method invocation:

```
int[] array = { 1, 2, 3 };  
int sum = sumNumbers(array, "odd");
```

the result should be 4. Nevertheless, due to a bug in the code, the result is 3.

Thanks to AD, with only this information (without knowing anything about the source code) we can identify the buggy method. For instance, if we debug this program

with the Declarative Debugger for Java (DDJ),¹ we obtain the following debugging session (questions are generated by DDJ, and answers are provided by the user):

```
Starting Debugging Session:
(1) sumNumbers({1,2,3},"odd")=3? No
(2) isEven(2)=true? Yes
(3) isEven(3)=false? Yes
(4) add(0,3)=3? Yes
Bug found in method "sumNumbers".
```

Hence, an AD session is just a dialogue where the debugger asks questions and the user answers them. The Java code associated with this debugging session is depicted in Figure 1.

```
int add(int x, int y) {
    return x + y;
}

boolean isEven(int x) {
    if ((x % 2) == 0)
        return true;
    return false;
}

int sumNumbers(int[] arrays, String eo) {
    int sum = 0;
    for (int i = 1; i <= arrays.length; i++) {
        if (eo.equals("even") && isEven(arrays[i]))
            sum = add(sum, array[i]);
        if (eo.equals("odd") && !isEven(arrays[i]))
            sum = add(sum, array[i]);
    }
    return sum;
}
```

Fig. 1. Java program to sum even or odd numbers in an array

What the debugger internally does, is to generate a data structure that represents the execution of the program. This data structure, often called Execution Tree (ET), is depicted in Figure 2. The ET has a node for each method invocation.² Each node normally contains a reference to the method that is being executed, the value of its arguments, the old and new values of the variables that may be changed within the execution, and its returned value. The debugger just traverses the ET asking the user about the validity of the nodes (i.e., nodes are marked as correct or wrong) until a buggy node is found. A node is buggy when it is wrong, and all of its children (if any) are correct.

The main properties of AD are the following:

Theorem 1 (Correctness of AD [23]). *Given an ET with a buggy node n , the method associated with n contains a bug.*

Theorem 2 (Completeness of AD [27]). *Given an ET with a bug symptom (i.e., the root is wrong), provided that all the questions generated by the debugger are correctly answered, then, a bug will eventually be found.*

¹ <http://www.dsic.upv.es/~jsilva/DDJ/>

² In the ET, nodes represent computations. Hence, depending on the underlying paradigm, they can represent methods, functions, procedures, clauses, etc. Our discussions in this paper can be applied to any paradigm, but, for the sake of concreteness, we will center the discussion on the object-oriented paradigm and our examples on Java.




Fig. 2. ET generated for the program in Figure 1

1.1 Motivation

We have been actively working in the area of AD for the last 10 years. This paper somehow summarizes and criticizes our own work to make a step forward. We claim that almost all current algorithmic debuggers—at least all that we know, including the most extended, which we compared in [7], and including our own implementations—have fundamental problems that were somehow inherited from the original formulation of AD [27].

In particular, we claim that the original formulation of AD, and most of the later definitions and implementations are obsolete with respect to the last advances in the practical side of AD. For instance, two important problems of standard definitions of AD are the granularity and the static nature of the errors found (AD reports a whole routine as buggy). We can illustrate these problems observing again the debugging session of Example 1: The whole method `sumNumbers` is pointed as buggy. This is very imprecise specially if `sumNumbers` was a big method with a lot of code. However, AD researchers and implementors are used to this behavior, and they could answer that this is the normal output of any algorithmic debugger. However, from an engineering perspective, this is quite surprising because the analysis performed by the debugger is by definition dynamic (in fact, the whole program is actually executed). Thus, the debugger should know that lines 36-39 are never executed, and thus they should not be reported as buggy. This lead us to our first proposition: The information reported by an algorithmic debugger should be dynamic and not static. That is, the output should be the part of a method that has been actually executed to produce the bug, instead of the whole method.

We think that this problem comes from the first implementations of AD and it has been inherited in latter theoretical and practical developments. In fact, if we execute this program with the debuggers: Buddha [26], DDT [4], Freja [21], Hat-Delta (and its predecessor Hat-Detect) [8], B.i.O. [3], Mercury’s Algorithmic Debugger [19], Münster Curry Debugger [18], Nude [20], DDJ [13], and HDJ [11], they all would output the whole `sumNumbers` as buggy.

Traditionally, AD reports a whole method as buggy. To reduce the granularity of errors, new techniques have appeared (see, e.g., [16, 5]) that allow us to debug inside a method. Unfortunately, the standard definition of ET is not prepared for that. In fact, some of the recent transformations defined for AD do not fit in the traditional definition of the data structures used in this discipline. For instance, the *Tree Balancing* technique presented in [15], or the *zooming* technique presented in [5] cannot be represented with standard AD data structures such as the *Evaluation Dependence Tree* [24].

This lack of a common theoretical framework with standard data structures that are powerful enough as to represent recent developments makes researchers to reinvent

the wheel once and again. In particular, we have observed that researchers (including ourselves) have produced local and partial formalizations to define their systems for a particular language and/or implementation. These theoretical developments are hardly reusable in other languages, and thus, they only serve as a formal description of their system, or as a means to prove results.

1.2 Contributions

In this work, we propose a new redefinition of AD in such a way that:

1. It is paradigm- and language-independent, and thus it is reusable by other researchers.
2. It is a conservative generalization of the traditional formulation of AD, in such a way that many previous AD techniques are a particular case of this new formulation.
3. It is formulated in a way that definitions of the data structures, properties, strategies, and algorithms are specified separately, so that they can be reused and/or concretized to a particular case.
4. It states that the output of an algorithmic debugger should contain dynamic information (i.e., it should not include non-executed code).
5. It allows the debugger to ask questions about the code inside a method (and not only about the whole methods).

We do not claim that our definitions cannot be further generalized, but we do claim that this generalization is appropriate and enough for the current state of the art (contrarily to the current formulation). We also provide some UML models to clarify and categorize different ideas introduced in the recent bibliography, so that AD techniques can be classified with them. Finally, we provide properties that researchers should report about their AD techniques, because these properties allow other researchers and developers to know whether a given technique is useful for a particular application or not.

1.3 Outline of the Paper

The rest of the paper is structured as follows: First, in Section 2 we discuss the related work. Then, in Section 3 we present our new formulation of AD. It is structured in four parts: the ET, search strategies for AD, transformations for AD, and an AD scheme. Finally, Section 4 concludes.

2 Related Work

Algorithmic debugging has been applied to all mature languages. All current implementations use a sort of ET to represent computations. Even in those lazy implementations of AD where the execution of the front-end and the back-end is interleaved (see, e.g., [22]), the construction of the ET is needed before debugging. Along the years each paradigm has adopted a well-defined and studied data structure to implement the ET.

2.1 A Little Bit of History

The first notion of ET appeared in the seminal work by Shapiro [27] in the context of the logic paradigm. Shapiro used refutation trees as ETs. Later implementations of AD in the logic paradigm such as NU-Prolog [29, 2] also used refutation trees.

In the context of the functional paradigm, the data structure used was proposed by Henrik Nilsson and Jan Sparud: The Evaluation Dependence Tree (EDT). They first proposed this data structure as a record of the execution [24], and then, as an appropriate ET for AD [25]. The EDT is particularly useful to represent lazy computations by hiding non-computed terms. In fact, the own EDT can be computed lazily as in [22]. The most successful implementations of AD for the functional paradigm are based on the EDT. Notable cases are the debuggers Freja [21], Hat-Delta [8], and Buddha [26].

In multi-paradigm languages such as Mercury, TOY, or Curry, the ET is also a proof tree, or an EDT. Examples of debuggers for these languages are the Mercury Debugger [19], the Münster Curry Debugger [18], DDT [4], and B.i.O [3].

In the imperative paradigm, a redefinition of the EDT was used. It has been often called *Execution Tree* [10, 13], but, conceptually, it is equivalent to the EDT, and it can be seen as a dynamic version of the *Call Graph* where every single call uses a different node in the graph, and thus no cycles are possible (i.e., it is a tree).

2.2 Modern Implementations

All the debuggers mentioned in previous sections are somehow “standard” in the sense that they are based on the standard definition of ET (either refutation trees or the EDT). However, in the last 5 years, there has been a new trend in AD tools: They implement new techniques that go beyond the standard definition of ET. Contrarily to the previously described tools, modern algorithmic debuggers are not standalone tools. They are plugins that can be integrated as part of an IDE. Examples of these debuggers are JHyde [12] and HDJ [11], being both of them part of Eclipse. Precisely because they are integrated into a development environment, they have direct access to dynamic information—they can even manipulate the JVM at runtime—that can be used to enhance the debugging session. In particular, the following techniques go beyond standard ETs: (i) *Tree compression* hides nodes of the ET (it breaks the standard parent-child relation in the ET). (ii) *Tree balancing* introduces new artificial nodes in the ET (it breaks the standard definition of ET node). (iii) *Loop expansion* and (iv) *ET zooming* decompose ET nodes (they break the standard definition of ET node). We are not aware of any definition of ET able to represent the previous four techniques.

3 Paradigm-Independent Redefinition of Algorithmic Debugging

Some of our last developments for AD cannot be formalized with the standard AD formulation. In a few cases, we just skipped the formalization of our technique and provided an implementation. In other cases we wanted to prove some properties, and thus we formalized (for one specific language, e.g., Java) the part of the system affected by those properties. Other developments were done for other paradigms, e.g., the functional paradigm, and we also formalized a different part of the system with different

data structures. We have observed other researchers doing the same, and clearly, this is due to the lack of a standard solution.

We want to provide a definition of AD that is paradigm-independent (i.e., it can be used by either imperative- or declarative- languages). From the best of our knowledge, there does not exist such a formal definition of algorithmic debugging. Hence, in this section we formulate algorithmic debugging in an abstract way. The main generalization of our new formulation is considering that ET nodes are not necessarily routines as in previous definitions (see, e.g., [24]). Contrarily, we allow ET nodes to contain any piece of code. This permits AD to report any code as buggy, and not only routines, thus potentially reducing the granularity of errors to single expressions.

In the following, we will only call *Execution Tree* to our new definition, and we will call *Routine Tree* (RT) to the traditional definition (that we also formalize in the next sections). Because our new definition is a conservative generalization, the RT is a particular case of the ET as it can be observed in the UML model of Figure 3.




Fig. 3. Debugging Hierarchy

Observe that an execution node can be specialized depending on the piece of code it represents. In particular, we specialize three kinds of execution nodes named *Routine Node*, *Projection Node*, and *Collapse Node*. They correspond to definitions already existing in the literature (see [10, 15]), but other kinds of nodes could appear in the future.

3.1 The Execution Tree

In this section we introduce some notation and formalize the notion of Execution Tree used in the rest of the paper. We want to keep the discussion and definitions in this section paradigm-independent. Hence, we consider programs as state transition systems.

Definition 1 (Program). A program $P = \{W, I, R\}$ consists of:

- W : A set of states.
- I : A set of starting states, such that $I \subseteq W$.
- R : A transition relation, such that $R \subseteq W \times W$.

Definition 2 (Computation). A computation is a maximal sequence of states s_1, s_2, \dots such that:

- s_1 is a starting state, i.e., $s_1 \in I$.
- $(s_i, s_{i+1}) \in R$ for all $i \geq 1$ (and $i \leq n - 1$, if the sequence is of the finite length n).

A finite segment s_i, s_{i+1}, \dots, s_j where $1 \leq i < j \leq n$ is called a subcomputation.

In the source code of a program, we consider statements³ as the basic execution unit. Therefore, in the following, the source code of a program P is a set of statements st_1, st_2, \dots, st_n that produces the computation s_0, s_1, \dots, s_m for a given starting state s_0 . We also use the notion of *code fragment* of a program P , which refers to any subset of statements in the source code of P that produces a subcomputation s_i, \dots, s_j with $0 \leq i < j \leq m$. Code fragments often represent functions or loops in a program, but they can also represent blocks, single statements, or even function calls together with the whole called function.

We assume that each state in W is composed of pairs variable-value. The initial and final states, s_i and s_j , describe the effects of a given code fragment c . All three together form a *code behavior*.

Definition 3 (Code Behavior). Given a code fragment c and a computation $C_c = s_0, \dots, s_n$ of c for a given initial state s_0 , the code behavior of C_c is a triple $(s_0, \mathcal{P}_{C_c}(c), s_n)$, where $\mathcal{P}_{C_c}(c)$ is the projection of c modulo C_c that contains the minimum subset of c needed to produce the computation s_0, \dots, s_n .

Code behavior corresponds to the questions asked by the debugger. These questions are along the lines of *Should the code c with the initial state s_0 produce the final state s_n ?*, or *Code c produced s_n from s_0 , is that correct?*. Many previous definitions of AD defined the code behavior as the triple (s_0, c, s_n) , which corresponded to the execution of a routine c , and usually the debugger only needed to show the call to c instead of showing both the call to c and the own routine c . This definition, however, introduces two important novelties:

- It allows c to be any code fragment, and not only a routine.
- It substitutes c by a projection of c modulo C_c , thus the code associated to a code behavior only contains the code needed to execute that behavior.

This dynamic notion is much more precise than the usual static notion that considers (the complete code of) a routine. For the sake of clarity, we sometimes use (s_0, d, s_n) meaning $(s_0, \mathcal{P}_{C_c}(c), s_n)$ when talking about a generic code behavior.

Definition 4 (Oracle Answers Set). Given a code fragment c , and the universe \mathcal{U} of all possible subcomputations of c (for all possible initial states), the oracle answers set \mathcal{A} of c is a set that contains one tuple (\mathcal{B}, a) for each subcomputation $C_c \in \mathcal{U}$ where:

1. \mathcal{B} is the code behavior of C_c , and
2. a is a boolean.

³ Note the careful use of the word “statement” to refer to either imperative instructions, declarative expressions, etc.

Intuitively, the oracle answers set represents the gold standard that determines whether any code behavior is correct. It is used as a reference point against which one can compare computations to determine whether they are correct or wrong.

We are now in a position to define the nodes of an execution tree.

Definition 5 (Execution Node). Let C_c be a computation of a code fragment c . In addition, let \mathcal{A} be the oracle answers set of c . The execution node associated with C_c is a pair (\mathcal{B}, S) where:

1. \mathcal{B} is the code behavior of C_c and
2. S is the state of the node, which can be either:
 - undefined, or
 - the correctness of C_c with respect to \mathcal{A} : $\begin{cases} \text{correct} & \text{if } (C_c, \text{true}) \in \mathcal{A} \\ \text{wrong} & \text{if } (C_c, \text{false}) \in \mathcal{A} \end{cases}$

Observe that an execution node contains (inside \mathcal{B}) the source code $\mathcal{P}_{C_c}(c)$ responsible of the computation it represents. Hence, if this node is eventually declared as buggy, its associated code is uniquely identified. Observe also that this code determines whether the node is a routine node or another kind of node depending on the code fragment that it represents.

In order to properly define execution trees, we need to define first a relation between execution nodes that specifies the parent-child relation.

Definition 6 (Execution Nodes Dependency). Let N be a set of execution nodes. Given a computation C_c with associated execution node $n_c \in N$, and a subcomputation $C_{c'}$ of C_c with associated execution node $n_{c'} \in N$, we say that n_c directly depends on $n_{c'}$ (expressed as $n_c \xrightarrow{N} n_{c'}$) if and only if there does not exist a subcomputation $C_{c''}$ of C_c with associated execution node $n_{c''} \in N$ such that $C_{c'}$ is a subcomputation of $C_{c''}$.

Observe that this dependency relation is intransitive, which is needed to define the parent-child relation in a tree. Hence, provided that we have three execution nodes, n_1, n_2, n_3 , if $n_1 \xrightarrow{N} n_2 \xrightarrow{N} n_3$ then $n_1 \not\xrightarrow{N} n_3$. Finally, we define an *execution tree*. It essentially represents the execution of a code in a structured way where each node represents a sub-execution of its parent. Formally,

Definition 7 (Execution Tree). Let C_c be a computation of a code fragment c . An Execution Tree (ET) of C_c is a tree $T = (N, E)$ where:

- $\forall n \in N$, n is the execution node associated with a subcomputation of C_c ,
- The root of the ET is the execution node associated with C_c ,
- $\forall (n_1, n_2) \in E$. $n_1 \xrightarrow{N} n_2$.

This definition is a generalization of the usual call tree (CT), which in turn comes from the refutation trees initially defined for AD in [27]. One important difference between them is that, given a computation C_c of a code fragment c , the CT associated with C_c is unique because it is only formed of routine nodes. In contrast, there exist different valid ETs associated with C_c due to the flexibility introduced by execution nodes (i.e., with routine nodes only one set N is possible, while with execution nodes different sets N

are possible). This relation 1 to n is interesting because it leaves room for transforming the ET and still being valid. Contrarily, the CT cannot be transformed because it would stop being a valid CT.

Once the ET is built, the debugger traverses the ET asking the oracle about the correctness of the information stored in each node. Using the answers, the debugger identifies a *buggy node* that is associated with a *buggy code* of the program.

Definition 8 (Buggy Node). Let $T = (N, E)$ be an execution tree. A buggy node of T is an execution node $n = (\mathcal{B}, \mathcal{S}) \in N$ where:

- (i) $\mathcal{S} = \text{wrong}$, and
- (ii) $\forall n' = (\mathcal{B}', \mathcal{S}') \in N, (n, n') \in E . \mathcal{S}' = \text{correct}$.

Moreover, we say that a buggy node n is traceable if and only if:

- (iii) $\forall n' = (\mathcal{B}', \mathcal{S}') \in N, (n', n) \in E^* . \mathcal{S}' = \text{wrong}$.

We use E^* to refer to the symmetric and transitive closure of E . This is the usual definition of buggy node (see, e.g., [23]): a wrong node with all its children correct. We also introduce the notion of *traceable*. Roughly, traceable buggy nodes are those buggy nodes that can be directly responsible of the wrong behavior of the program (their effects are visible in the root of the tree), or at least, there is a path of wrong nodes from the root (which must be wrong) to them. This property makes them debuggable by all AD strategies that are variants of Top-Down (see [28]).

Lemma 1 (Buggy Code). Let T be an ET with a buggy node $((s, d, s'), \mathcal{S})$ whose children are $((s_1, d_1, s'_1), \mathcal{S}_1), ((s_2, d_2, s'_2), \mathcal{S}_2) \dots ((s_n, d_n, s'_n), \mathcal{S}_n)$. Then $d \setminus \bigcup_{1 \leq i \leq n} d_i$ contains a bug.

Proof (Buggy Code). Trivial adaptation from the proof by Lloyd [17] for Prolog.

3.2 Routine Tree

In this section we formalize the notion of RT used in most AD literature as a particular case of the ET. We call *routine tree* to this specialization of the ET to make explicit its multi-paradigm nature, because routines can refer to functions, procedures, methods, predicates, etc. We first define a *routine node*, which is a specialization of an execution node.

Definition 9 (Routine Node). A routine node is an execution node $((s_0, \mathcal{P}_{C_c}(c), s_n), \mathcal{S})$ where code fragment c only contains:

- a routine call r , together with
- all the code of the routines directly or indirectly called from r .

Therefore, in a routine node, s_0 and s_n are, respectively, the states just before and after the execution of the routine called. Almost all implementations reduce c to the routine call, and they skip the code of the own routine.

Definition 10 (Routine Tree). A routine tree is an execution tree where all nodes are routine nodes.

3.3 Search Strategies for AD

Once the ET is built, AD uses a search strategy to select one node. During many years, the main goal of most AD researchers has been the definition of better strategies to reduce the search space after every answer, and to reduce the complexity of the questions. A survey of search strategies for AD can be found in [28]. In our formalization, a search strategy is just a function that analyzes the ET and returns an execution node (either the next node to ask, or the buggy node).

Definition 11 (Strategy). A search strategy is a function whose input is an execution tree $T = (N, E)$ and whose output is an execution node $n = (\mathcal{B}, \mathcal{S}) \in N$ such that:

1. $\mathcal{S} = \text{undefined}$, or
2. n is a buggy node.

3.4 AD Transformations

Some of the last research developments in AD have been focussed on the definition of transformations of the ET. The goal of these transformations is to improve its structure before the debugging session so that search strategies become more efficient. Some of these transformations cannot be represented with a routine tree. For this reason, we want to include this section to classify the kinds of transformations that have been defined so far, and establish a hierarchy so that future transformations can be also classified.

There exist three essential elements in the front-end of an algorithmic debugger. The modification of any of them can lead to a different final output of the front-end (i.e., a different ET). Therefore, we classify the transformations in three different levels:

- Transformations of the *source code*: Transformations of the source code such as *inlining* are used to reduce the size of the ET by hiding routines. Contrarily, transformations such as *loops to recursion* [14] are used to augment the size of the ET to reduce the granularity of the buggy code reported (a loop instead of a routine). In general, users should not be aware of the internal transformations applied by the debugger, thus the code fragment shown to the user should be the original code.
- Transformations of the *execution*: Transforming the way in which the source code is executed can change the ET generated. One example is changing eager evaluation by lazy evaluation. Another example is passing arguments by value instead of passing them by reference. We are not aware of any implementation that includes transformations of this kind.
- Transformations of the *ET*: Transforming the ET can significantly reduce the number of questions generated. In general, the ET is transformed with the aim of making search strategies to behave as a dichotomic search. Hence, they try to produce balanced ETs [15], or also deep trees that can be cut in the middle. Other transformations such as *Tree compression* [9] try to avoid the repetition of questions about the same routine.

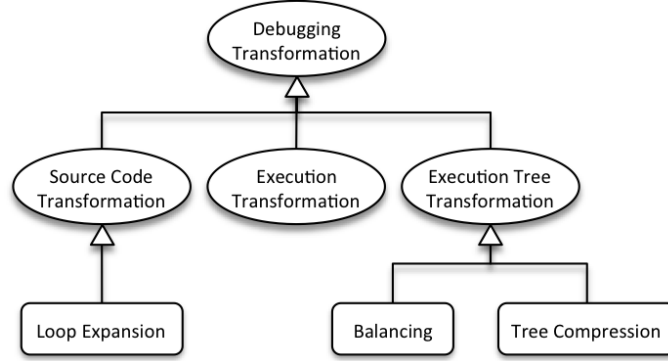


Fig. 4. Transformations Hierarchy

Given a computation C_c of a code fragment c , being $T = (N, E)$ the execution tree of C_c , we represent with $\mathcal{T}(T)$ either: (i) a transformation T' of T , (ii) the ET of a computation C'_c transformed from C_c , or (iii) the ET of a computation $C_{c'}$ of a code fragment c' transformed from c . Hence, $\mathcal{T}(T)$ is the ET produced after applying one or more of the previously described transformations. Note that we use T to refer to the ET that would have been generated without applying any transformation. A transformation must have some specific properties to be useful, otherwise, we have the risk of improving the structure of the ET at the cost of losing buggy nodes. We propose four different properties to measure the impact of a transformation on a given ET: *Buggy Node Completeness*, *Buggy Code Completeness*, *Buggy Code Reduction*, and *Bug Existence*. The last one must be satisfied by any transformation to ensure that AD is still able to find at least one bug. All of them are formally presented bellow.

Property 1 (Buggy Node Completeness). Given an ET $T = (N, E)$, and its transformed version $\mathcal{T}(T) = (N', E')$, $\forall n \in N$, if n is a buggy node in T then n is a buggy node in $\mathcal{T}(T)$.

Because all buggy nodes belong to both ETs, then, by Definition 7, the same buggy code behaviors are detectable. Hence, all bugs that can be localized in the ET generated before applying the transformation are still detectable after applying the transformation. Many transformations cannot satisfy Property 1, but they can satisfy a relaxed version:

Property 2 (Buggy Code Completeness). Given an ET $T = (N, E)$, and its transformed version $\mathcal{T}(T) = (N', E')$, $\forall n = ((s_0, d, s_n), \mathcal{S}) \in N$, n is a buggy node in T . $\exists n' = ((s'_0, d', s'_n) \in N'$, n' is a buggy node in $\mathcal{T}(T)$ and $d = d'$.

This property ensures that all the code that can be reported as buggy in the original ET, can be also reported as buggy in the transformed ET. However, the specific context that produced the bug is not necessarily the same in the original and in the transformed version. There exist transformations that violate the last condition ($d = d'$). However, this is not a problem provided that d' is a subset of d , because this means that the

code reported as buggy is smaller in the transformed code (i.e, precision is increased by reducing the granularity of the bug found).

Property 3 (Buggy Code Reduction). Given an ET $T = (N, E)$, and its transformed version $\mathcal{T}(T) = (N', E')$, $\forall n = ((s_0, d, s_n), S) \in N$, n is a buggy node in T . $\exists n' = ((s'_0, d', s'_n) \in N'$, n' is a buggy node in $\mathcal{T}(T)$ and $d \subseteq d'$.

Sometimes, these properties are further relaxed to a bug existence property.

Property 4 (Bug Existence). Given an ET $T = (N, E)$, and its transformed version $\mathcal{T}(T) = (N', E')$, If $\exists n \in N$, n is a buggy node in T , then $\exists n' \in N'$, n' is a buggy node in $\mathcal{T}(T)$.

Property 4 is the minimum exigible requirement to ensure that the transformation has not hidden all bugs. Thus, all transformations must satisfy Property 4.

In Figure 4 we classify three transformations in the state of the art. Two of them, *tree balancing* and *loop expansion* produce ETs that are not routine trees. The technique *tree balancing* satisfies Property 1, except for *collapse nodes* that only satisfy Property 2, thus, all buggy nodes (except those replaced by collapsed nodes) are preserved. However buggy nodes can stop being traceable after this transformation. The technique *loop expansion* satisfies Property 3, but not Property 2 (and thus, nor Property 1). The reduction in the buggy code is obtained introducing the ability in the ET of detecting bugs inside loops. The technique *tree compression* satisfies Property 2, but not Property 1.

3.5 An AD Scheme

Finally, we describe Algorithm 1, a general schema of an algorithmic debugger that includes all phases, from the generation of the ET to the bug reported. This algorithm gives an idea of how and when, the ET, the transformations, the oracle, and the search strategies participate in the whole debugging process.

The main function performs the two phases of AD (Lines 1-2) and then returns a buggy code of the program (Line 3). In the first phase (function *getExecutionTree*) the ET is created performing all possible transformations in the source code (Line 1), in the execution (Line 3) and in the ET (Line 5). Once the ET is created, the second phase (function *debugTree*) starts. During this phase, the debugger traverses the ET selecting nodes with a search strategy (Line 2). Function (*selectNode*) implements one of the search strategies in the literature. There has been a lot of research for more than a decade concerning which should be the node to ask. A survey can be found in [28]. No matter what strategy is used, *selectNode* returns a node to ask (the state of the node is *undefined*), or a buggy node (the state of the node is *wrong* (Line 3)). Once a node has been selected, the debugger asks the oracle about its correctness (Line 5). With the answer of the oracle, the debugger updates the state of the nodes of the ET (Lines 5-7). Note that the answer of the oracle can affect the state of several nodes. This can effectively change the structure of the ET, and thus, at this moment, a new ET transformation could be used to optimize the ET (Line 8). Then, the process is repeated selecting more nodes. When the strategy finds a buggy node (Lines 3-4) or it cannot select more nodes (Line 1) the second phase finishes and the debugger returns (see

Algorithm 1 Main algorithm of an Algorithmic Debugger

Input: A program P and its input i .

Output: A buggy code C in P , or \perp if no bug is detected in P .

Initializations: $\mathcal{A} = \emptyset$ // Set of answers provided by the oracle

begin

1) $T = \text{getExecutionTree}(P, i)$

2) $n = \text{debugTree}(T)$

3) **return** $\text{getCode}(n)$

end

function $\text{debugTree}(T = (N, E))$

begin

1) **while** $(\exists (\mathcal{B}', S') \in N, S' = \text{undef} \vee \text{wrong})$

2) $(\mathcal{B}, S) = \text{selectNode}(T)$ // Strategy

3) **if** $(S = \text{wrong})$ **then**

4) **return** (\mathcal{B}, S)

5) $\text{answer} = \text{askOracle}(\mathcal{B})$

6) $\mathcal{A} = \mathcal{A} \cup (\mathcal{B}, \text{answer})$

7) $\text{updateStates}(\mathcal{A}, N)$

8) $T = \text{executionTreeTransformations}(T)$

9) **return** \perp

end

function $\text{getExecutionTree}(P, i)$

begin

1) $P' = \text{sourceCodeTransformations}(P)$

2) $\mathcal{E}_{P'} = \text{executeProgram}(P', i)$

3) $\mathcal{E}'_{P'} = \text{executionTransformations}(\mathcal{E}_{P'})$

4) $T = \text{generateExecutionTree}(\mathcal{E}'_{P'})$

5) $T' = \text{executionTreeTransformations}(T)$

6) **return** T'

end

function $\text{getCode}(n)$

begin

1) **if** $(n = ((s_0, d, s_n), S))$ **then**

2) **return** d

3) **return** \perp

end

getCode function) the buggy code associated with the buggy node found, or it returns a message indicating that there does not exist a bug, respectively. The last case happens, e.g., when all nodes are reported as *correct* (it is indicated with \perp in Line 3).

4 Conclusion

In this paper we report about some of the problems identified in the current state of the art of AD. One of the problems identified is that much of the recent work in the area does not fit into the standard notions and definitions of AD. In particular, we claim that practically all current definitions of the ET are obsolete with respect to the new techniques proposed.

To solve this situation we propose a generalization of AD able to represent all existent AD transformations. We make this abstraction considering theoretical developments done for a particular language or technique that are generalized, but also considering novel implementations of AD that include techniques that have not been formalized. We also propose some properties useful to classify the recently proposed AD transformations, and also future developments.

The main objectives of this work are two: First, putting together different ideas that have appeared in many works of AD. Putting these ideas together provides a wide perspective that allows us to make a step forward in the abstraction and generalization of the theoretical side of AD. In addition, it allows us to provide classifications and taxonomies to help understanding the state of the art. Second, our new formulation of AD

tries to save time. Many researchers have defined once and again similar concepts used in different languages and tools. We provide a paradigm-independent definition that is general enough to represent all current techniques, and it can be easily instantiated to any particular language.

5 Acknowledgements

This work has been partially supported by the EU (FEDER) and the Spanish *Ministerio de Economía y Competitividad (Secretaría de Estado de Investigación, Desarrollo e Innovación)* under Grant TIN2013-44742-C4-1-R and by the *Generalitat Valenciana* under Grant PROMETEOII/2015/013. David Insa was partially supported by the Spanish Ministerio de Educación under FPU Grant AP2010-4415.

References

1. Eclipse. Available from URL: <http://www.eclipse.org/>, 2003.
2. T. Barbour and L. Naish. Declarative debugging of a logical-functional language. Technical report, University of Melbourne, 1994.
3. B. Brael and H. Siegel. Debugging lazy functional programs by asking the oracle. In O. Chitil, Z. Horváth, and V. Zsóka, editors, *Proceedings of the 19th International Symposium on Implementation and Application of Functional Languages*, volume 5083 of *Lecture Notes in Computer Science*, pages 183–200. Springer Berlin Heidelberg, 2007.
4. R. Caballero. A Declarative Debugger of Incorrect Answers for Constraint Functional-Logic Programs. In *Proceedings of the 2005 ACM-SIGPLAN Workshop on Curry and Functional Logic Programming (WCFLP'05)*, pages 8–13, New York, USA, 2005. ACM Press.
5. R. Caballero, E. Martin-Martin, A. Riesco, and S. Tamarit. EDD: A declarative debugger for sequential erlang programs. In E. Ábrahám and K. Havelund, editors, *20th International Conference Tools and Algorithms for the Construction and Analysis of Systems (TACAS 2014)*, volume 8413 of *Lecture Notes in Computer Science*, pages 581–586. Springer, apr 2014.
6. R. Caballero, A. Riesco, A. Verdejo, and N. Martí-Oliet. Simplifying questions in maude declarative debugger by transforming proof trees. In G. Vidal, editor, *21st International Symposium Logic-Based Program Synthesis and Transformation (LOPSTR 2011)*, volume 7225 of *Lecture Notes in Computer Science*, pages 73–89. Springer, jul 2011.
7. D. Cheda and J. Silva. State of the practice in algorithmic debugging. *Electronic Notes in Theoretical Computer Science*, 246:55–70, 2009.
8. T. Davie and O. Chitil. Hat-delta: One Right Does Make a Wrong. In A. Butterfield, editor, *Proceedings of the 17th International Workshop on Implementation and Application of Functional Languages (IFL'05)*, page 11, sep 2005.
9. T. Davie and O. Chitil. Hat-delta: One Right Does Make a Wrong. In *Proceedings of the 7th Symposium on Trends in Functional Programming (TFP'06)*, apr 2006.
10. P. Fritzson, N. Shahmehri, M. Kamkar, and T. Gyimóthy. Generalized Algorithmic Debugging and Testing. *ACM Letters on Programming Languages and Systems (LOPLAS)*, 1(4):303–322, dec 1992.
11. J. González, D. Insa, and J. Silva. A New Hybrid Debugging Architecture for Eclipse. In *Proceedings of the 23th International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR 2013)*, volume 8901 of *Lecture Notes in Computer Science (LNCS)*, pages 183–201. Springer, 2014.

12. C. Hermanns and H. Kuchen. Hybrid Debugging of Java Programs. In M. J. Escalona, J. Cordeiro, and B. Shishkov, editors, *Software and Data Technologies*, volume 303 of *Communications in Computer and Information Science*, pages 91–107. Springer Berlin Heidelberg, 2013.
13. D. Insa and J. Silva. An Algorithmic Debugger for Java. In *Proceedings of the 26th IEEE International Conference on Software Maintenance (ICSM 2010)*, pages 1–6, 2010.
14. D. Insa and J. Silva. Automatic transformation of iterative loops into recursive methods. *Information and Software Technology*, 58:95–109, 2015.
15. D. Insa, J. Silva, and A. Riesco. Speeding Up Algorithmic Debugging Using Balanced Execution Trees. In M. Veanes and L. Viganò, editors, *Proceedings of the 7th International Conference Tests and Proofs (TAP 2013)*, volume 7942 of *Lecture Notes in Computer Science (LNCS)*, pages 133–151. Springer, jun 2013.
16. D. Insa, J. Silva, and C. Tomás. Enhancing Declarative Debugging with Loop Expansion and Tree Compression. In E. Albert, editor, *Proceedings of the 22th International Symposium on Logic-based Program Synthesis and Transformation (LOPSTR 2012)*, volume 7844 of *Lecture Notes in Computer Science (LNCS)*, pages 71–88. Springer, sep 2012.
17. J. Lloyd. Declarative Error Diagnosis. *New Generation Computing*, 5(2):133–154, 1987.
18. W. Lux. Münster Curry User’s Guide. Available at URL: <http://danae.uni-muenster.de/~lux/curry/user.pdf>, may 2006.
19. I. D. MacLarty. *Practical Declarative Debugging of Mercury Programs*. PhD thesis, University of Melbourne, 2005.
20. L. Naish, P. W. Dart, and J. Zobel. The NU-Prolog Debugging Environment. In A. Porto, editor, *Proceedings of the 6th International Conference on Logic Programming (ICLP’89)*, pages 521–536, Lisboa, Portugal, 1989.
21. H. Nilsson. *Declarative Debugging for Lazy Functional Languages*. PhD thesis, Linköping, Sweden, may 1998.
22. H. Nilsson. How to look busy while being as lazy as ever: the implementation of a lazy functional debugger. *Journal of Functional Programming*, 11(6):629–671, 2001.
23. H. Nilsson and P. Fritzon. Algorithmic Debugging for Lazy Functional Languages. *Journal of Functional Programming*, 4(3):337–370, 1994.
24. H. Nilsson and J. Sparud. The evaluation dependence tree: an execution record for lazy functional debugging. Technical report, Department of Computer and Information Science, Linköping, 1996.
25. H. Nilsson and J. Sparud. The evaluation dependence tree as a basis for lazy functional debugging. *Automated Software Engineering*, 4(2):121–150, 1997.
26. B. Pope. *A Declarative Debugger for Haskell*. PhD thesis, The University of Melbourne, Australia, 2006.
27. E. Shapiro. *Algorithmic Program Debugging*. MIT Press, 1982.
28. J. Silva. A Survey on Algorithmic Debugging Strategies. *Advances in Engineering Software*, 42(11):976–991, nov 2011.
29. B. Thompson and L. Naish. A guide to the nu-prolog debugging environment. Technical report, University of Melbourne, 1997.