

Moldable Exceptions

Andrei Chiş
feenk gmbh
Wabern, Switzerland
andrei.chis@feenk.com

Tudor Gîrba
feenk gmbh
Wabern, Switzerland
tudor.girba@feenk.com

Oscar Nierstrasz
feenk gmbh
Wabern, Switzerland
oscar.nierstrasz@feenk.com

Abstract

Debugging is hard. Interactive debuggers are mostly the same. They show you a stack, a way to sample the state of the stack, and, if the debugger is live, a way to step through execution. The standard interactive debugger for a general-purpose programming language provided by a mainstream IDE mostly offers a low-level interface in terms of generic language constructs to track down and fix bugs. A custom debugger, such as those developed for specific application domains, offers alternative interfaces more suitable to the specific execution context of the program being debugged. Custom debuggers offering contextual debugging views and actions can greatly improve our ability to reason about the current problem. Implementing such custom debuggers, however, is non-trivial, and poses a barrier to improving the debugging experience. In this paper we introduce *moldable exceptions*, a lightweight mechanism to adapt a debugger’s interface based on contextual information provided by a raised exception. We present, through a series of examples, how moldable exceptions can enhance a live programming environment.

CCS Concepts: • Software and its engineering → Software maintenance tools.

Keywords: Exceptions, debuggers, customization.

ACM Reference Format:

Andrei Chiş, Tudor Gîrba, and Oscar Nierstrasz. 2024. Moldable Exceptions. In *Proceedings of the 2024 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software (Onward! ’24)*, October 23–25, 2024, Pasadena, CA, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3689492.3690044>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
Onward! ’24, October 23–25, 2024, Pasadena, CA, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1215-9/24/10

<https://doi.org/10.1145/3689492.3690044>

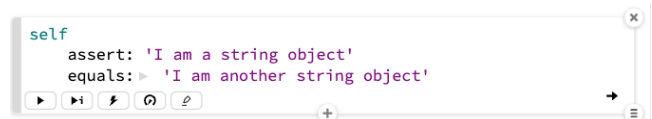
1 Introduction

Debuggers are unloved beasts. They provide a generic and low-level interface to explore and debug the run-time state of a running program. Many debugging problems can be better tackled by understanding the nature of the exception that was raised and caused the debugger to be activated.

There have been numerous efforts to develop custom debuggers for various application domains and domain-specific languages. These custom debuggers provide dedicated views and actions that are tailored to a specific application context. For example, an *object-centric debugger* [25] offers views and interactions that focus on specific objects rather than the run-time stack. Building a custom debugger is, however, a non-trivial task, so this does not happen too often. An *extensible debugger* (such as *deet* [13]) is designed so that it can be easily extended with new graphical views and debugging operations, but these extensions still represent a significant development effort. A *moldable debugger* [7] is a special kind of extensible debugger that can activate alternative debugger interfaces depending on the current execution context, however the development of these alternative debuggers is still non-trivial.

We propose a new, lightweight mechanism, called *moldable exceptions*, to dynamically adapt a moldable debugger using contextual information provided by the exception itself. In modern, object-oriented software, it is common practice to define dedicated classes of exceptions to signal individual run-time issues. Each exception therefore implicitly carries knowledge about the kind of issue being raised. For example in a parsing framework an exception can indicate that a given input cannot be parsed. Moldable exceptions leverage this knowledge by associating simple views and actions to be activated by a moldable debugger when that exception is raised.

Consider the following example. In [Figure 1](#) we see an assertion that compares two strings.¹ In a normal setting, this



```
self
assert: 'I am a string object'
equals: 'I am another string object'
```

Figure 1. A failing string comparison assertion. assertion will fail, yielding a standard debugger view, such as the one we see in [Figure 2](#). This is the standard debugger of

¹All the examples are written in Pharo Smalltalk (<https://pharo.org>), running in the open-source Glamorous Toolkit (GT) IDE (<https://gtoolkit.com>).

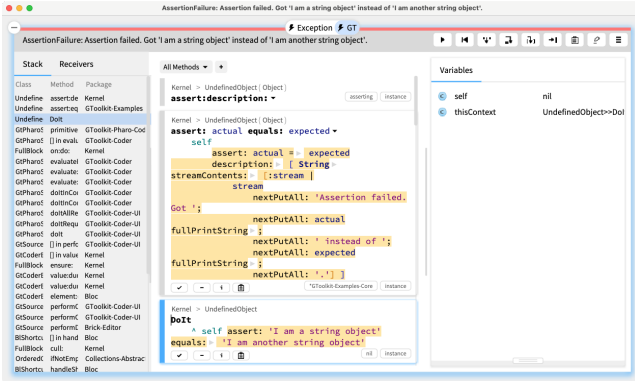


Figure 2. A generic debugger view.

the GT IDE, which resembles other generic debuggers in that it shows an error message at the top, a compact view of the reified run-time stack at the left, source code of each stack frame in the middle, an inspector on the variables of the current stack frame at the right, and a dashboard of buttons at the top right for stepping through the running program. It will typically take a developer some time to putter around in the debugger interface to understand the specific error (the strings do not match), and *why* the strings do not match.

Suppose that instead of seeing the generic debugger, we are offered a view that highlights the actual differences, as in Figure 3. Such a view not only homes directly in on the spe-

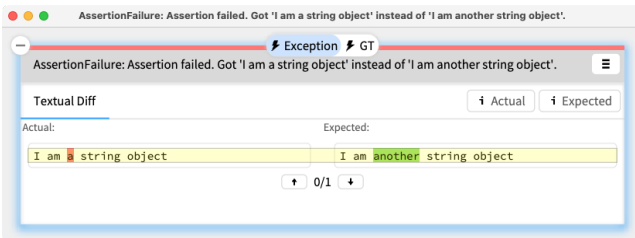


Figure 3. A string diff debugger view.

cific problem, but also highlights the individual differences in a dedicated “diff” view. Furthermore, since the diff view already exists as a component used in other applications, the development effort is close to zero. While in this introductory example, the difference is easy to spot, in other cases such as those presented in section 6, a diff view is essential to quickly identifying differences.

Moldable exceptions work as follows: when an exception is raised, the exception (an object) is caught and passed to the debugger. Every exception not only provides the debugger with the context it needs to generate the debugging UI, but it can also offer alternative views and actions. In the case of our implementation, this is achieved by the exception class providing specially annotated debugger extension methods.

This simple mechanism allows moldable exceptions to do three things:

- (i) provide domain-specific debugging views and actions,
- (ii) offer new debugger GUI interactions, and

- (iii) enable automated fixes (code and data transformations) for common programming errors.

We will illustrate these three points with several examples. In section 2 we introduce the *moldable inspector*, which allows objects to define custom views and actions when they are inspected. In section 3 we show how custom debugger views and actions can easily be defined in a similar way by adding simple annotated view and action methods to the class of a moldable exception. In section 4 we show how a richer debugger user interface can be provided in much the same way by leveraging existing GUI frameworks. We show an example of moldable exceptions enabling automated fixes in section 5. In section 6 we describe a case study in which we used moldable exceptions to improve a development workflow. We discuss the implementation details in section 7. We summarize our contributions and discuss some possible future work in section 8, and we discuss related work in section 9. We conclude in section 10.

2 The Moldable Inspector

Understanding moldable exceptions can be easier if we first introduce another example of a similar approach. Moldable exceptions build on the idea of a *moldable tool* [8], an IDE tool that adapts its behavior to a specific run-time application context. An example of such a tool is a *moldable object inspector* [9]. When a moldable object inspector is opened on an object it looks for any view or action methods defined in the class of that object, and uses those methods to create custom views.

Let us consider a typical example. Suppose we have an implementation of a Ludo² game. Players alternate in throwing a die and moving a token until one player reaches their goal square. In Figure 4 we see a classical *Raw* view supported by a typical object inspector. It just shows the state of the

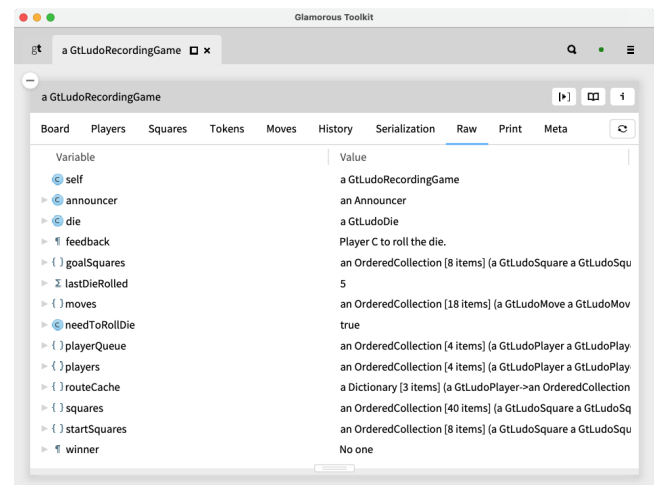


Figure 4. A classical “raw” view of a Ludo game.

²A simple game in which players move tokens around a board based on the roll of a die. <https://en.wikipedia.org/wiki/Ludo>

object as a list of instance variables and their values. In [Figure 5](#), however, we see an alternative, graphical *Board* view showing the current state of the game as a user would see it.

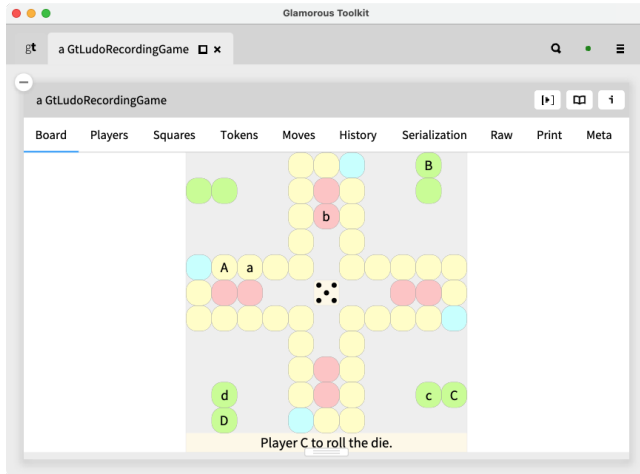


Figure 5. An interactive GUI view of a Ludo game.

In general, a variety of views may be more useful than the classical “raw” view, and this is also true for the Ludo game. In [Figure 6](#) we see a *Moves* view that lists all the moves of the game played thus far. By clicking on a move (at

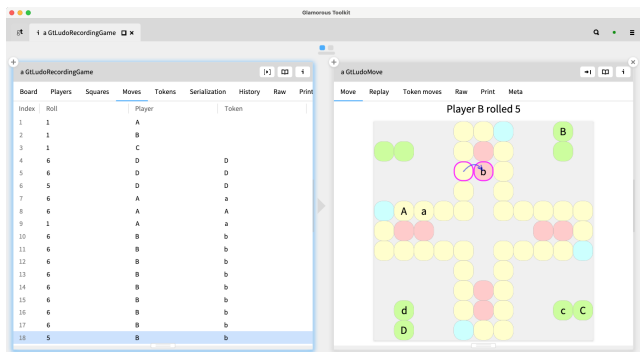



Figure 6. A custom historical *Moves* view of a Ludo game, and a custom graphical view of a selected move.

the left) you can then inspect (at the right) the corresponding `GtLudoMove` object, with its custom views, in this case showing the details of what happened in this move.

In [Figure 7](#) we see the source code defining the *Moves* view. This method uses just a few lines of code to create a browsable “columned list” of past moves with several columns for the details of each move.

In addition to views, objects can also define *actions* that encapsulate useful operations. For example, the game object in [Figure 5](#) defines an *Autoplay one move* action (with the icon ) that can be used to simulate a player’s turn.

View and action methods are recognized by the moldable inspector through dedicated annotations, in exactly the same way that a test runner tool in a classical IDE recognizes Java

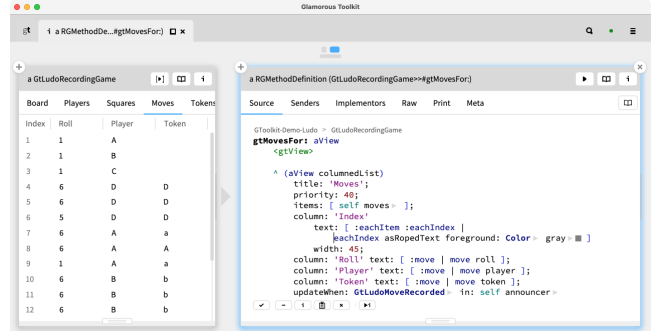


Figure 7. The *Moves* view next to its source code.

test case methods because they are tagged with a `@Test` annotation. In this case the annotation is `<gtView>` for views (and `<gtAction>` for actions), seen in the second line of the method defining the view.

3 Adding Custom Debugger Views

Moldable exceptions provide custom debugger views and actions in essentially the same way as the moldable inspector. Moldable exceptions are instances of an *Exception* class that has been extended with a dedicated method for each custom debugger view or action. These methods are annotated with a `<gtExceptionView>` pragma for views and a `<gtExceptionAction>` pragma for actions to extend the moldable debugger (i.e., as opposed to `<gtView>` and `<gtAction>` which extend the moldable inspector).

Let us consider that the Ludo game has been implemented with the help of *Design by Contract* [18]. Rolling a die when a player should move, or vice versa, constitutes a precondition violation, which raises a `LudoMoveAssertionFailure`. Similarly, if an attempt is made to move the wrong player’s token, this will trigger a precondition failure. Normally, this would fire up the classical debugger, as seen in [Figure 8](#). Although

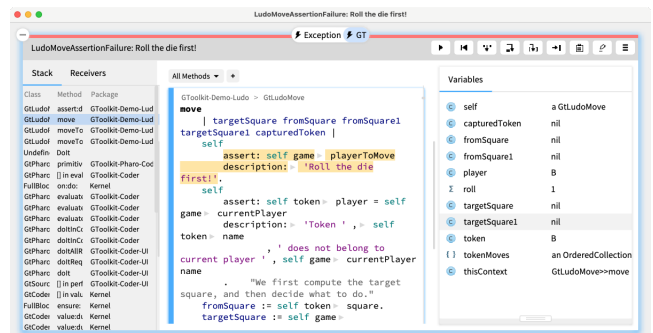


Figure 8. A classical debugger for a precondition failure. the precondition violation is clearly reported, the debugger interface is not ideal for tracking down the actual reason for the violation.

What we would perhaps like to see instead is the current state of the game, in addition to a history of the past moves. We could possibly find these by navigating through the existing debugger views, but why not show them directly? After

all, we know whenever this exception is raised, what it is that we would like to see. Furthermore, if we already have such views defined elsewhere (we do!), it is not a question of defining new views, but of reusing them in the context of the debugger. We just need to define two new view methods in the class `LudoMoveAssertionFailure`.

Here is the definition of the first view, which simply forwards (delegates) the view to another existing one. Specifically, the moldable debugger game view for a `LudoMoveAssertionFailure` simply reuses the moldable inspector view called `gtPositionsFor`: that is already defined in the game object.

```
1 gtGameViewFor: aView
2 <gtExceptionView>
3 ↑ aView forward
4 title: 'Game';
5 priority: 10;
6 object: [ move game ];
7 view: #gtPositionsFor:
```

Let us step through the code: (i) `gtGameViewFor:` is the name of the view method, which takes as its argument `aView`, the view to be defined. (ii) The method is annotated with `<gtExceptionView>`, which tells the moldable debugger to enable the view whenever `LudoMoveAssertionFailure` is raised. (iii) We return (`↑`) the result of sending forward to `aView`. (iv) We give the view a title and (v) a priority (the order in which views are shown – we set it to 10 rather than 1 to leave room for more important views later). (vi) We specify the object to forward to (the move’s game). (vii) We name the already existing view method to forward to (`gtPositionsFor`).

We similarly define a method for the history of past moves, and now the debugger, instead of showing us the classical debugger, will offer us the two views seen in [Figure 9](#). The Game view shows us the current game state graphically, and the Moves view shows us a browsable list of past moves. Note that we can always switch to the standard debugger by selecting the *GT* button at the top.

The *Moves* debugger view is also defined by reusing the existing object inspector *Moves* view:

```
1 gtMovesViewFor: aView
2 <gtExceptionView>
3 ↑ aView forward
4 title: 'Moves';
5 priority: 20;
6 object: [ move game ];
7 view: #gtMovesFor:
```

The assertion diff debugger view we saw earlier in [Figure 3](#) is similarly defined as a method of `AssertionFailure`.

```
1 AssertionFailure>>gtComparableTypesTextualDiffFor: aView
2 <gtExceptionView>
3 | assertionContext |
4 self gtHasStack ifFalse: [ ↑ aView empty ].
```



Figure 9. Custom debugger views for the Ludo game.

```
5 assertionContext := self
6 gtLocateAssertEqualsContextWithComparableTypes.
7 assertionContext ifNil: [ ↑ aView empty ].
8 ↑ aView forward
9 title: 'Textual Diff';
10 priority: 0;
11 object: [ assertionContext ];
12 view: #gtComparableTypesTextualDiffFor:
```

The key difference is that not every `AssertionFailure` is raised as the result of a comparison. For this reason, in line 6, the view will be suppressed (`↑ aView empty`) in case the assertion did *not* fail in the context of an `assert:equals:` check.

4 Building Domain-Specific Debugger Views

The examples we have seen so far have reused existing inspector views, but sometimes there is a need to develop a new kind of debugger view. This also need not necessarily imply a heavy implementation effort.

The *GT Scripter* is a tool used to “script” GUI interactions, mainly for testing purposes. Let us consider the case of assertion errors being raised while testing a user interface interaction with the help of the *Scripter*. The scenario is the following: (i) We create a notebook page called “Page One”. (ii) We add a text snippet with a link to the `Object` class. (iii) We assert the existence of the page. (iv) We click the link. (v) We check that a new page is created with a code browser for the `Object` class. In case something goes wrong, we would like the debugger to provide us with high-level views of the state of the *Scripter*.

We consequently introduce three dedicated debugger views. The *Scripter preview* (Figure 10) shows the result of the scripted interaction: (i) the notebook page has been created, containing a link to the class Object, and (ii) the link to the class has been clicked, opening a source code browser on the class.

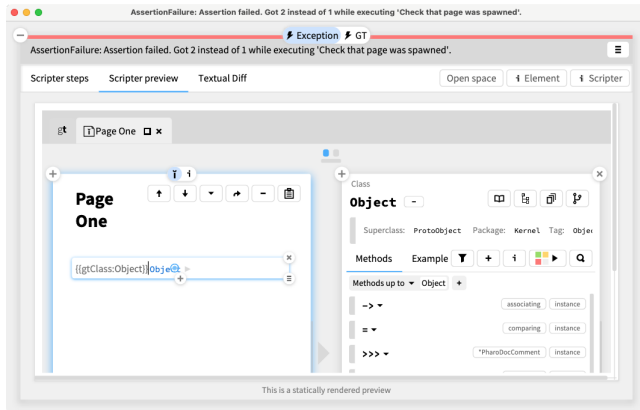


Figure 10. Scripter preview.

The *Scripter steps* view (Figure 11) shows us a graphical tree view of the steps performed by the GUI Scripter, as well as the assertions that have been checked. (In this case the tree has only two levels, but in general there may be many levels.) The green steps and assertions have succeeded, whereas the

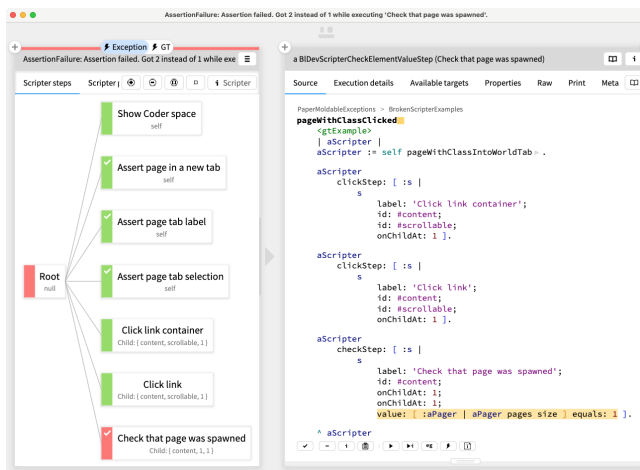


Figure 11. Scripter steps view.

red ones have failed. By clicking on any step or assertion node, we can see the corresponding code highlighted in the scripter method at the right. Here we see that the step *Check that page was spawned* has failed, and at the right we see the corresponding failure highlighted in the code.

Finally, the *Textual Diff* view (Figure 12) is the same one we have seen earlier, reused again. It just tells us that the check should verify that there are now two pages in the pager, not just one, as we can also see in the *Scripter preview*. It is the check specification that is at fault, not the GUI Interaction we are testing.

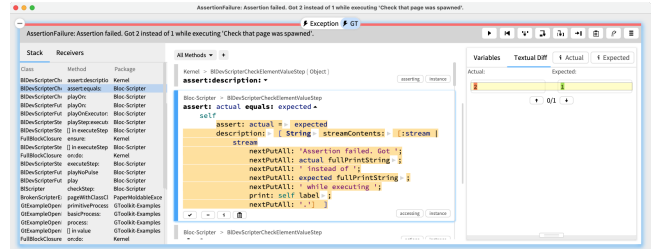


Figure 12. Scripter textual diff view.

However this time we show the Textual Diff view both in the new debugger interface created for this exception (Textual Diff in Figure 10), as well as an additional view within the standard GT debugger (Textual Diff tab in Figure 12). Hence, a moldable exception can also provide views and actions for the standard GT debugger.

Why do we need these views? The problem with the standard debugger is that, due to the way the GUI Scripter schedules the steps, the offending method (i.e., the pageWithClassClicked method we see in Figure 11) is not on the stack at the point where the exception is raised. Although it is possible to get at the information we seek, doing so is clumsy since we would have to navigate to the exception instance itself, and explore its internal state, and the classical stack view only confuses matters instead of helping us to debug the problem. This can be the case with many application domains, especially those that depend on event scheduling. The run-time stack does not do a good job of telling us how the events have been triggered, so another kind of view is needed.

In general, custom debugger views exist to save you the trouble of rummaging around the stack to find the information you need, and they can offer dedicated custom actions to fix the problem.

How hard is it to implement a domain-specific debugger interface? Extending a standard debugger with such views can be challenging. Typically such views would have to be added in a subclass of the standard debugger. With moldable exceptions, views only have to be added to the exception class itself. The *Scripter steps* view is implemented using a version of *Mondrian* [1, 19], a builder for graph-based visualizations. In Figure 13 we can see the entire source code of the view expanded in place, implemented in just four methods.³ The debugging view at the top just adds a “Inspect Scripter” button to the next method, an object inspector view for a Scripter. This in turn just delegates to the *Steps* view of a *BIDevScripterStep* object. Finally, this method embeds a *Mondrian* visualization implemented in 12 lines of code in the *paintTreeWith:* method.

Obviously this does not prove that all domain-specific debugger interfaces will be so tiny, but it does demonstrate

³NB: The figure is intended just to emphasize the size of the source code, not to be an invitation to try to read it.

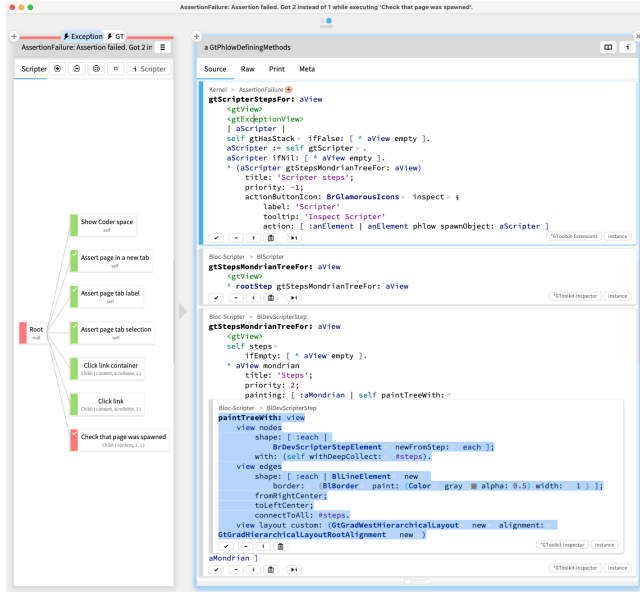


Figure 13. The Scriptor steps source code.

that a useful custom debugger view can be implemented in an almost trivial amount of code.

5 Enabling Automated Fixes

Many cases of common programming errors can be automatically repaired. Consider the case providing to an API an object of the wrong type, or of the right type but the wrong state. While the first kind of mistake could arguably be caught by a static type system, detecting objects being in a wrong state is rather a run-time issue, typically caught by a precondition. Some common cases can be fixed by rewriting the client code. In Figure 14 we see the source code for a custom inspector view that incorrectly returns a view instead of a view empty in the preamble. This custom debugger view



Figure 14. Catching an “empty view” error.

decorates the source code with an explanation pointing out the likely errors, and links (in yellow) to the offending source code.

Since such preambles are a common idiom in defining GT inspector views, and the error is also not uncommon, it becomes easy to fix with the help of a transformation.

Figure 15 shows the result of performing the *Fix & retry* action: a refactoring widget is generated that proposes a code transformation. Clicking on the *Apply refactoring and proceed* button will transform the source code, leave the debugger, and evaluate the rewritten code. Given how GT inspector views are executed, in this particular case, after applying the transformation that fixes the code of the view, it is enough to restart the top call frame. In other cases it might be needed to walk over the call stack to find the right place from where to restart execution.

A custom debugger action, such as *Fix & retry* or *Retry* is defined in the same way as a custom debugger view: it is a method with a particular annotation, in this case being `<gtExceptionAction>` rather than `<gtExceptionView>`, but the mechanism is essentially the same. We see the code of this action in the listing below. Since this is a debugging action we need to be able to control the run-time execution. To make this possible, methods creating debugging views and actions can take as a second parameter a context object that gives them access to the current run-time execution.

```

1  gtDebugActionFixItFor: anAction inContext: aContext
2  <gtExceptionAction>
3  (self shouldApplyInContext: aContext )
4  ifFalse: [ ↑ anAction noAction ].
5  ↑ anAction dropdown
6  label: 'Fix & retry';
7  icon: BrGlamorousVectorIcons play ;
8  priority: 50;
9  id: #'gtDebugAction-fixit';
10 preferredExtent: 650@350;
11 content: [ self refactoringElementWithAction: [
12     aContext debuggingSession restartTopContext.
13     aContext debugger resumeAndCloseDebugger ]]

```

The default behavior of such fixits is to open a debugger with the possibility of applying a proposed code transformation. If, on the other hand, such code transformations should be applied automatically, this can be configured. In our proof-of-concept implementation of moldable exceptions, this is done by evaluation the following code, which sets a flag in a globally accessible Singleton.

```

1  GtMoldableExceptionTransformationsSettings
   defaultInstance allowAutomaticTransformations.

```

The mechanism for managing unhandled exceptions can then optionally consult this flag to decide whether to spawn a debugger or apply the fix (section 7 gives more details).

As before, the custom debugger views and actions are mostly built from existing components, so the implementation effort is low.

6 A Case Study of Moldable Exceptions

As a case study, we applied moldable exceptions within a company that develops applications for managing insurance



Figure 15. Transforming an empty view error.

policies, to improve the development process by integrating custom comparison views into the debugger.

One part of the application consists in generating various documents related to insurance contracts. The creation of documents is verified using tests that check that a recreated document matches an expected version that has been previously saved. Whenever a comparison fails, a developer needs to assess the differences between the two documents, and then, in case those differences are justified, update the saved version with the recreated one, or in case they are not expected, track down and fix the root cause of the difference.

In the established workflow, a developer uses a dedicated tool to compare the recreated and the expected versions, and in case the difference is acceptable, update the expected document with the new one. In case the difference indicates a bug, the developer has to run the test again using the debugger to explore the live execution. With this workflow, developers have to choose the tool *before* they run the test, which means that they typically need to run a test twice. This can be a problem as some tests take minutes to run.

With moldable exceptions we integrated custom debugger views that show the differences between the recreated and the expected versions, and provide custom actions to update the expected document with the new one and resume executing the test, directly in the debugger. To achieve this we rely on the fact that tests raise `ComparisonFailure` exceptions when two documents do not match.

Figure 16 shows the debugger opened as a result of a failed comparison of two PDF documents. In this case the exception provides multiple diff views: one at the image level that visually highlights the parts that differ between the two documents (Figure 16), a second one at the level of the text contained by the PDF documents (Figure 17), and another at the level of the raw document format (Figure 18). Apart from the diff views, the exception also provides an *Accept*

action that replaces the saved expected document with the new one, and continues the execution of the test.

The *Image Diff* view and the *Accept* action both reuse existing logic, and in a few lines of code integrate it into the debugger. This integration unifies the flow of comparing two documents and exploring why the failure happens using a debugger. Developers can now simply run tests, and when a comparison fails choose how to proceed.

7 Implementation Details

The key to implementing moldable exceptions is the mechanism used to inform a moldable debugger of the relevant views and actions it provides. In our implementation every exception can provide one or more *debugger specifications*. Exceptions define these specifications through methods annotated with `<gtDebuggerSpecification>`. Below we see the method in the `Exception` class, at the top of the exception class hierarchy, that defines a custom debugger for every exception.

```

1 Exception>>gtExceptionDebuggerSpecification
2 <gtDebuggerSpecification>
3 ↑ GtMoldableExceptionSpecificDebuggerSpecification
4   forException: self
    
```

A debugger specification contains all the information needed by a debugger to locate actions and views. Figure 19 shows the default properties defined by such a specification.

By default, as mentioned in section 3, the moldable GT debugger looks for methods annotated with a `<gtExceptionView>` pragma. An exception can change the pragma if needed, or provide a list of other pragmas. The default specification also provides several other properties, including:

- Debugging Targets: objects reachable from the exception where to look for actions and views.



Figure 16. A view showing visual differences between two PDF documents (differences are highlighted in magenta).




Figure 17. A view showing textual differences between the text content of two PDF documents.




Figure 18. A view showing textual differences between document at the level of the PDF format.

- Debugging Actions: pragmas for debugging actions.
- Title, Icon: for displaying a debugger selector.
- Priority: a number used to determine the order in which alternative debuggers are offered.
- Activation Predicate: a condition to decide whether to show the debugger or not based on the exception state; by default the debugger is shown if any debugging views are found.
- Available Automatically: if the debugger is active, indicates whether it can be shown by default.
- Debugger Class: The class of the debugger that is created from this specification.

By default the debugger created from a specification searches for debugging views and actions in the class of the exception. However, through the Debugging Targets property an

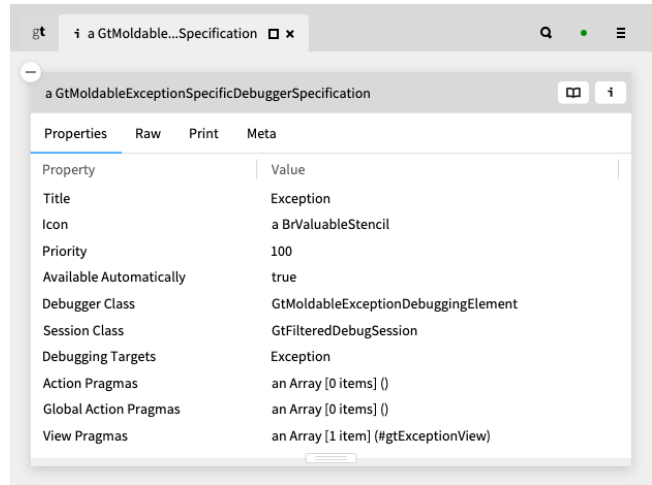
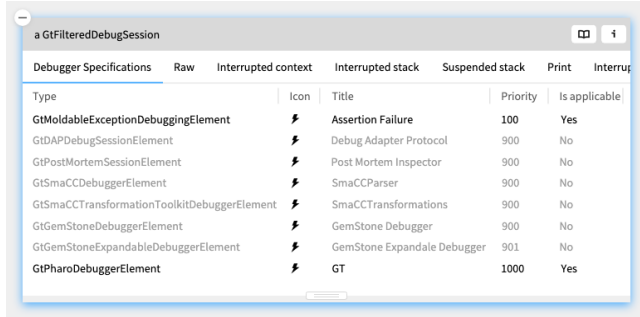


Figure 19. Default properties for a debugging specification.

exception can change that, and provide any other list of objects. That makes it possible for other objects in the system reachable from the exception objects to provide debugging actions and views.

To summarize, when an exception is raised, the moldable debugger looks for debugger specifications in exception methods annotated with <gtDebuggerSpecification> and uses them to create custom debuggers. These custom debuggers are then shown when their activation predicate is true. For example, Figure 20 shows the list of debugger specifications located when the assertion failure from Figure 3 is raised. We see that in total eight specifications are found, but just the one for the default debugger together with the one provided by the exception are active, and result in a custom debugger view being shown.

To enable moldable exceptions to perform automated transformations instead of opening the debugger, we provide a



Type	Icon	Title	Priority	Is applicable
GtMoldableExceptionDebuggingElement	⚡	Assertion Failure	100	Yes
GtDAPDebugSessionElement	⚡	Debug Adapter Protocol	900	No
GtPostMortemSessionElement	⚡	Post Mortem Inspector	900	No
GtSmaCCDebuggerElement	⚡	SmaCCParser	900	No
GtSmaCCTransformationToolkitDebuggerElement	⚡	SmaCCTransformations	900	No
GtGemStoneDebuggerElement	⚡	GemStone Debugger	900	No
GtGemStoneExpandableDebuggerElement	⚡	GemStone Expandable Debugger	901	No
GtPharoDebuggerElement	⚡	GT	1000	Yes

Figure 20. A list of debugger specifications found when an exception is raised.

dedicated trait [10] that exceptions can use. This trait provides an implementation for the signal method shown below. This method is executed whenever an exception is raised. By default in case an exception is not handled in the system a debugger is opened. This override checks if automatic transformations are allowed for the current exception (line 4), and if they are, asks the exception for its transformation (lines 5) and executes it if there is anything to transform (lines 6-7). Transformations can apply code changes, modify the state of objects, or perform other operations, such as deleting no longer needed files.

```

1 TgtMoldableExceptionSignalWithTransformation>>signal
2 <debuggerCompleteToSender>
3 | transformation |
4 self canApplyAutomaticTransformation iffFalse: [↑super
5   signal].
5 transformation := self
6   gtMoldableExceptionTransformation.
6 transformation shouldTransform iffFalse: [↑super signal].
7 transformation performTransformation

```

8 Discussion and Future Directions

Moldable exceptions are just objects augmented with annotated methods that create alternative views and actions when they are used to spawn a moldable debugger. The only changes needed to an application are: (i) if there is not already a dedicated exception class for the exception to be molded, introduce one, and (ii) add dedicated exception view methods to that class to mold the debugger.

The previous extension mechanism for the moldable debugger required developers to extend the debugger by creating subclasses [7]. Although the old framework considerably reduced the cost of building a custom debugger, they still required a non-trivial implementation effort, consisting of several hundreds of lines of code ([7], table 7). While both the moldable inspector and the (old) moldable debugger were available in GT up to now, currently GT comes with 3200 views and 280 actions with an average of 12 lines of code, but only 8 custom debuggers (such as the one in Figure 21).

This indicates that the extension mechanism plays a significant role in developers extending the IDE, and that the cost of creating views and actions can be very low. Due to this, through moldable exceptions we aim to apply the same extension mechanism present in the moldable inspector for creating custom debuggers. In the new approach, a moldable debugger simply needs to recognize whether a raised exception has custom views defined (in our case by detected annotated methods of its class), and then activate those views in its user interface.

At the time of writing we have implemented some 26 custom debugger views, in an average of under 12 lines of code. Over half of these are forwarding views that reuse (delegate to) existing views previously defined as object inspector views. Several more are simple text, list or tree views, and only three create custom graphical widgets.

The simplest debugging views are just like inspector views, and many of our examples just forward (delegate) to existing views, but more generally exceptions have access to the reified run-time stack at the time that they are raised, so debugger views can extract and present arbitrary run-time information. The *Scripter steps* view we saw earlier in Figure 11 offers an example. The same approach could be used, for example, to present or highlight just the “interesting” stack frames, for example in an event-driven application, just the methods that are responsible for processing events.

At this time, moldable exceptions only offer the possibility to provide alternative views and actions, but not alternative ways to step through the execution. Currently this is only possible in our setting by providing a completely separate debugger implementation. In Figure 21 we see such a dedicated debugger for SmaCC [5], the Smalltalk compiler compiler framework that has been spawned on an invalid fragment of Java code. The SmaCC debugger offers the possibility to

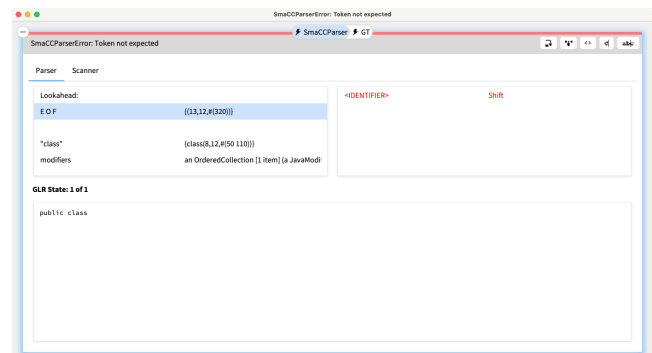


Figure 21. A custom SmaCC debugger on an invalid Java snippet.

step through the grammar rules of a parser and explore its execution state.

Our proof-of-concept implementation of moldable exceptions presented here heavily leverages the existing infrastructure for moldable inspector views in GT, but in principle there is nothing to prevent its application to other programming languages and IDEs. The basic idea is simple: any moldable tool must be prepared, whenever it is created, to examine the execution context of the objects it is initialized with, and use that context to adapt its behavior. In the case of a moldable debugger, this context is provided by the exception raised. Custom debugger views and actions must then be provided by the specific exception raised. One way to provide such views and actions is through specially annotated methods, but other means could be used, such as naming conventions, or a registry of debugger extensions. The precise extension method will depend on the language technology available.

9 Related Work

Early debuggers provided a command-line interface to inspect and step through the execution of a running program or a *post mortem* “core dump.” With the advent of GUI-based IDEs, debuggers were also updated to offer interactive views to explore the execution state of a program and to step through its execution [26].

Studies of developer debugging behavior have shown that many developers have difficulty using debuggers and often shy away from them. McCauley *et al.* report: “We note, however, that our experiences with debugging tools is that many of them use execution traces as a method of assisting students to understand the execution of a program. Tubaishat [27] characterized the use of execution traces as an example of a shallow reasoning technique” [16]. And Beller *et al.* note that “Debuggers are difficult to use. Another reason given by interviewees, even though seasoned engineers, was that ‘the debugger is a complicated beast’ (I2) and that ‘debuggers that are available now are certainly not friendly tools and they don’t lend toward self-exploration’” [2].

An early example of an extensible debugger is *deet* [13]. Extensions to *deet* were written Tcl/Tk or the Korn shell. The focus of the extension mechanism in *deet* is on adding new debugger features, rather than providing debugging support for specific application contexts.

The Java Virtual Machine Tools Interface (JVMTI) [21] is an API for tools to inspect the state and control the execution of Java applications. Agents can be written in any native language that supports C calling conventions. Custom views must be built from scratch and integrated in your environment.

The VSCode Debugger Adapter Protocol (DAP) [20] makes it possible to implement a generic debugger for a development tool that can communicate with different debuggers via Debug Adapters.

mbeddr [28] is a collection of languages and language extensions, mainly focusing on C, built with the JetBrains MPS language workbench [6]. *mbeddr* includes an extensible debugger framework [22] that focuses on providing support for the individual language extensions, not application specific customizations.

Bousse *et al.* have developed a generic omniscient (“back in time”) debugger that can be adapted to different domain-specific languages (DSLs) with moderate effort [3]. They use DAP to provide new controls for DSLs in VSCode [12]. Their work focuses more on the stepping semantics than on providing custom views.

There have been numerous efforts to develop specialized debuggers for DSLs using a variety of approaches. Wu *et al.* report on grammar-driven generation of DSL debuggers [29]. Lindeman *et al.* leverage the Spoofox language workbench [14] to declaratively define DSL debuggers [15]. D2X [4] provides an API for DSLs to define debugger extensions. All of these DSL debugger approaches focus on providing general debugging for a given DSL, not for more finely-grained application debugging issues.

Babylonian programming [23] is an approach in which live examples are integrated into the code editor, and offers various specialized views [24] of the run-time behavior of the examples. The current implementation of Babylonian/S does not offer a means to define application-specific views for examples.

Sindarin is an extensible Pharo-based debugger that offers a rich API for scripting debugger extensions [11]. The use case for Sindarin is rather different from that of moldable exceptions, allowing developers to enter debugger scripts during a debugging session, rather than as a way to enable specialized debugger behavior in response to specific kinds of exceptions. It would be interesting to explore the use of Sindarin to implement custom stepping behavior.

The Moldable Debugger [7] is an example of a moldable tool [8] that adapts its behavior to a specific run-time application context. The original moldable debugger, however, still required some significant implementation effort to define an alternative debugger interface, much like the other approaches we have cited from literature. Similarly, the original moldable inspector [9] did not offer a very easy mechanism to provide custom inspector views depending on the object being inspected. The adoption of the extension mechanism of using annotated methods to define inspector customizations led, over time, to the development of a large number of custom inspector views in the current version of GT (over 3000 at last count).

Disciplined use of exceptions is a well-established practice in object-oriented software development [17]. The key insight and original contribution of moldable exceptions is to leverage exceptions as the hook to molding the behavior of the debugger. Although it is common practice to add textual descriptions to exceptions, for example, by overriding

the `toString()` method in Java, associating debugging views to exceptions is novel. No previous approach to extensible debuggers, including the original moldable debugger, has taken this approach.

10 Conclusion

Interactive debuggers are all basically the same. Command-line debuggers offer commands to sample the current execution state and step through the running code. Graphical debuggers that show us the run-time stack, and offer buttons instead of commands to step through the code, but they are still all the same. The trouble with this is that every debugging problem is different, but debuggers all show us the same thing, and they offer the same kinds of actions to step through the code.

Our key contribution is to show how a debugger can be dynamically molded with custom views and actions that depend on the specific exception raised. The mechanism is lightweight, and many useful extensions can be defined with minimal implementation effort.

Acknowledgments

Many thanks to Alexandre Bergel, Steven Costiou, and Timo Kehrer for their reviews of a draft of this paper.

References

- [1] Vanessa Peña Araya, Alexandre Bergel, Damien Cassou, Stéphane Ducasse, and Jannik Laval. 2013. Scripting Visualizations with Mondrian. In *Deep Into Pharo*. Square Bracket Associates, 20. <https://inria.hal.science/hal-00858725/document>
- [2] Moritz Beller, Niels Spruit, Diomidis Spinellis, and Andy Zaidman. 2018. On the dichotomy of debugging behavior among programmers. In *Proceedings of the 40th International Conference on Software Engineering*, 572–583. <https://doi.org/10.1145/3180155.3180175>
- [3] Erwan Bousse, Dorian Leroy, Benoit Combemale, Manuel Wimmer, and Benoit Baudry. 2018. Omniscient debugging for executable DSLs. *Journal of Systems and Software* 137 (2018), 261–288. <https://doi.org/10.1016/j.jss.2017.11.025>
- [4] Ajay Brahmakshatriya and Saman Amarasinghe. 2023. D2X: An extensible conteXtual Debugger for Modern DSLs. In *Proceedings of the 21st ACM/IEEE International Symposium on Code Generation and Optimization*, 162–172. <http://groups.csail.mit.edu/commit/papers/2023/ajay-cgo23-d2x.pdf>
- [5] John Brant, Jason Lecerf, Thierry Goubier, Stéphane Ducasse, and Andrew Black. 2017. Smacc: a compiler-compiler. <http://books.pharo.org/booklet-Smacc/pdf/2018-10-21-Smacc-Compiler.pdf>
- [6] Fabien Campagne. 2014. *The MPS language workbench: volume I*. Vol. 1. Fabien Campagne.
- [7] Andrei Chiş, Marcus Denker, Tudor Girba, and Oscar Nierstrasz. 2015. Practical domain-specific debuggers using the Moldable Debugger framework. *Computer Languages, Systems & Structures* 44, Part A (2015), 89–113. <https://doi.org/10.1016/j.cl.2015.08.005> Special issue on the 6th and 7th International Conference on Software Language Engineering (SLE 2013 and SLE 2014).
- [8] Andrei Chiş, Tudor Girba, Juraj Kubelka, Oscar Nierstrasz, Stefan Reichhart, and Aliaksei Syrel. 2017. Moldable Tools for Object-oriented Development. In *PAUSE: Present And Ulterior Software Engineering*, Bertrand Meyer Manuel Mazzara (Ed.). Springer, Cham, 77–101. https://doi.org/10.1007/978-3-319-67425-4_6
- [9] Andrei Chiş, Tudor Girba, Oscar Nierstrasz, and Aliaksei Syrel. 2015. The Moldable Inspector. In *Proceedings of the 2015 ACM International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Pittsburgh, PA, USA) (*Onward! 2015*). ACM, New York, NY, USA, 44–60. <https://doi.org/10.1145/2814228.2814234>
- [10] Stéphane Ducasse, Oscar Nierstrasz, Nathanael Schärli, Roel Wuyts, and Andrew P. Black. 2006. Traits: A Mechanism for fine-grained Reuse. *TOPLAS: ACM Transactions on Programming Languages and Systems* 28, 2 (March 2006), 331–388. <https://doi.org/10.1145/1119479.1119483>
- [11] Thomas Dupriez, Guillermo Polito, Steven Costiou, Vincent Aranega, and Stéphane Ducasse. 2019. Sindarin: A versatile scripting API for the Pharo debugger. In *Proceedings of the 15th ACM SIGPLAN International Symposium on Dynamic Languages*, 67–79. <https://doi.org/10.1145/3359619.3359745>
- [12] Josselin Enet, Erwan Bousse, Massimo Tisi, and Gerson Sunyé. 2023. Protocol-Based Interactive Debugging for Domain-Specific Languages. *The Journal of Object Technology* 22, 2 (2023). <https://doi.org/10.5381/jot.2023.22.2.a6>
- [13] David R. Hanson and Jeffrey L. Korn. 1997. A Simple and Extensible Graphical Debugger. In *IN WINTER 1997 USENIX CONFERENCE*, 173–184. https://www.usenix.org/legacy/publications/library/proceedings/ana97/full_papers/hanson/hanson.html
- [14] Lennart C. L. Kats and Eelco Visser. 2010. The Spoofox Language Workbench. Rules for Declarative Specification of Languages and IDEs. In *OOPSLA'10: Proceedings of the 25th International Conference on Object-Oriented Programming, Systems, Languages, and Applications*, Martin Rinard (Ed.). Reno/Tahoe, NV, USA, 444–463. <https://doi.org/10.1145/1869459.1869497>
- [15] Ricky T. Lindeman, Lennart C.L. Kats, and Eelco Visser. 2011. Declaratively Defining Domain-specific Language Debuggers. *SIGPLAN Not.* 47, 3 (Oct. 2011), 127–136. <https://doi.org/10.1145/2189751.2047885>
- [16] Renee McCauley, Sue Fitzgerald, Gary Lewandowski, Laurie Murphy, Beth Simon, Lynda Thomas, and Carol Zander. 2008. Debugging: a review of the literature from an educational perspective. *Computer Science Education* 18, 2 (2008), 67–92. <https://doi.org/10.1080/08993400802114581>
- [17] Bertrand Meyer. 1988. *Disciplined Exceptions*. TR-EI-22/EX. Interactive Software Engineering, Goleta, CA. <https://se.inf.ethz.ch/old/people/meyer/publications/methodology/exceptions.pdf>
- [18] Bertrand Meyer. 1992. Applying Design by Contract. *IEEE Computer (Special Issue on Inheritance & Classification)* 25, 10 (Oct. 1992), 40–52. <https://doi.org/10.1109/2.161279>
- [19] Michael Meyer, Tudor Girba, and Mircea Lungu. 2006. Mondrian: An Agile Visualization Framework. In *ACM Symposium on Software Visualization (SoftVis'06)*. ACM Press, New York, NY, USA, 135–144. <https://doi.org/10.1145/1148493.1148513>
- [20] Microsoft. 2021. Debug Adapter Protocol. <https://microsoft.github.io/debug-adapter-protocol>
- [21] Oracle. 2024. JVM Tool Interface version 22.0. <https://docs.oracle.com/en/java/javase/22/docs/specs/jvmti.html>
- [22] Domenik Pavletic, Markus Voelter, SyedAoun Raza, Bernd Kolb, and Timo Kehrer. 2015. Extensible Debugger Framework for Extensible Languages. In *Reliable Software Technologies – Ada–Europe 2015*, Juan Antonio de la Puente and Tullio Vardanega (Eds.). Lecture Notes in Computer Science, Vol. 9111. Springer International Publishing, 33–49. https://doi.org/10.1007/978-3-319-19584-1_3
- [23] David Rauch, Patrick Rein, Stefan Ramson, Jens Lincke, and Robert Hirschfeld. 2019. Babylonian-style Programming – Design and Implementation of an Integration of Live Examples Into General-purpose Source Code. *Art Sci. Eng. Program.* 3, 3 (2019), 9. <https://doi.org/10.22152/PROGRAMMING-JOURNAL.ORG/2019/3/9>
- [24] Patrick Rein, Christian Flach, Stefan Ramson, Eva Krebs, and Robert Hirschfeld. 2024. Broadening the View of Live Programmers: Integrating a Cross-Cutting Perspective on Run-Time Behavior into a Live

- Programming Environment. *Art Sci. Eng. Program.* 8, 3 (2024), 38. <https://doi.org/10.22152/programming-journal.org/2024/8/13>
- [25] Jorge Ressia, Alexandre Bergel, and Oscar Nierstrasz. 2012. Object-Centric Debugging. In *Proceedings of the 34rd international conference on Software engineering (Zurich, Switzerland) (ICSE '12)*. <https://doi.org/10.1109/ICSE.2012.6227167>
- [26] Jonathan B Rosenberg. 1996. *How debuggers work: algorithms, data structures, and architecture*. John Wiley & Sons, Inc.
- [27] Abdallah Tubaishat. 2001. A Knowledge Base for Program Debugging. In *Proceedings of the ACS/IEEE International Conference on Computer Systems and Applications*. 321. <https://doi.org/10.1109/AICCSA.2001.934005>
- [28] Markus Voelter, Bernd Kolb, Tamás Szabó, Daniel Ratiu, and Arie van Deursen. 2017. Lessons learned from developing mbeddr: a case study in language engineering with MPS. *Software & Systems Modeling* (09 Jan. 2017). <https://doi.org/10.1007/s10270-016-0575-4>
- [29] Hui Wu, Jeff Gray, and Marjan Mernik. 2008. Grammar-driven Generation of Domain-specific Language Debuggers. *Softw. Pract. Exper.* 38, 10 (Aug. 2008), 1073–1103. <https://doi.org/10.1002/spe.v38:10>

Received 2024-04-25; accepted 2024-08-08