

# Accepted Manuscript

## $\mu$ -DSU: A Micro-Language Based Approach to Dynamic Software Updating

Walter Cazzola, Ruzanna Chitchyan, Awais Rashid, Albert Shaqiri

PII: S1477-8424(17)30041-6  
DOI: [10.1016/j.cl.2017.07.003](https://doi.org/10.1016/j.cl.2017.07.003)  
Reference: COMLAN 269

To appear in: *Computer Languages, Systems & Structures*

Received date: 14 March 2017  
Revised date: 3 July 2017  
Accepted date: 4 July 2017

Please cite this article as: Walter Cazzola, Ruzanna Chitchyan, Awais Rashid, Albert Shaqiri,  $\mu$ -DSU: A Micro-Language Based Approach to Dynamic Software Updating, *Computer Languages, Systems & Structures* (2017), doi: [10.1016/j.cl.2017.07.003](https://doi.org/10.1016/j.cl.2017.07.003)

This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.



**Highlights**

- This paper presents  $\mu$ -DSU a linguistic approach to dynamic software updating (DSU)
- The  $\mu$ -DSU framework is implemented on top of Neverlang
- $\mu$ -DSU exploits Neverlang's characteristic to permit a dynamic update of the interpreter during its execution
- $\mu$ -DSU has the benefit of easing the introduction and the update of nonfunctional concerns in a running application
- $\mu$ -DSU update process leaves the application code untouched reducing the complexity of the updating process and fostering adaptation reuse on a broader scale.
- Demonstration studies are provided

# $\mu$ -DSU: A Micro-Language Based Approach to Dynamic Software Updating

Walter Cazzola<sup>a,\*</sup>, Ruzanna Chitchyan<sup>c</sup>, Awais Rashid<sup>b</sup>, Albert Shaqiri<sup>a</sup>

<sup>a</sup>Computer Science Department, Università degli Studi di Milano, Italy.

<sup>b</sup>School of Computing and Communications, Lancaster University, United Kingdom.

<sup>c</sup>Department of Informatics, University of Leicester, United Kingdom.

---

## Abstract

Today software systems play a critical role in society's infrastructures and many are required to provide uninterrupted services in their constantly changing environments. As the problem domain and the operational context of such software changes, the software itself must be updated accordingly.

In this paper we propose to support dynamic software updating through language semantic adaptation; this is done through use of micro-languages that confine the effect of the introduced change to specific application features. Micro-languages provide a logical layer over a programming language and associate an application feature with the portion of the programming language used to implement it. Thus, they permit to update the application feature by updating the underlying programming constructs without affecting the behaviour of the other application features.

Such a linguistic approach provides the benefit of easy addition/removal of application features (with a special focus on non-functional features) to/from a running application by separating the implementation of the new feature from the original application, allowing for the application to remain unaware of any extensions. The feasibility of this approach is demonstrated with two studies; its benefits and drawbacks are also analysed.

---

## 1. Introduction

Today software applications govern many of our daily activities: from stock trading to traffic control, and in-body integrated insulin release pumps. Many such applications must provide continuous, uninterrupted services, and if interrupted, could lead to risk to life and safety (e.g., nuclear reactor control software) and substantial financial loss (e.g., software for stock trading).

Yet, as operational environment and expectations on such applications evolve, they do need to be maintained and updated. This problem is well-known, with several *dynamic software updating* (DSU) approaches proposed over the years to try to solve it, e.g., [1–3]. To date, evolution with DSUs is a two-steps process: i) the source code is modified according to the evolutionary needs and ii) the changed code is deployed on the running

---

\*Corresponding author, E-mail: cazzola@di.unimi.it

system using code injection, class reloading and dynamic/multiple method tables. Yet, these solutions inevitably result in long-term performance decay (e.g., JavAdaptor [1, 4] and DUSC [5]) or impose limitations on what can or cannot be updated (e.g., JRebel [2]).

In spite of all these efforts, software evolution is still work in progress, especially in supporting the application extension with new non-functional features<sup>1</sup> whose *a posteriori* design and addition is often cumbersome and could be jeopardized during application execution. In literature, we find many examples of this kind of evolution, e.g., for persistence [6, 7], parallelization [8, 9], optimisation [6] and pluggable type systems [10].

To fully understand the problem, consider, for example, a sequential application that could be sped up on a multi-core laptop computer by executing each independent stage of a loop<sup>2</sup> in parallel on a different core. This parallelization is clearly a non-functional feature that can fast the goal achievement but its absence does not compromise its feasibility. The parallelization can be supported by manually (or using an IDE) extracting the body of the loops with independent stages into methods and replacing these loops with the spawning of the corresponding method calls into the available cores. This update must be applied to all the loops which must be parallelised. Several new methods, which have no direct relation to the application logic, must be introduced into the code. This makes the code more obscure and complex. To complicate the scenario, let us consider the fact that we would like to pass from a sequential to a parallel version (which is more energy intensive but quicker) when our laptop is plugged into the mains, and revert back to sequential processing when the laptop runs on the battery. This is a clearly feasible change that most DSUs can support but it can hardly be automated and a continuous manual switch from one version to the other could lead to errors, e.g., forgetting to parallelise some of the loops. Also, though in theory it is feasible to keep two separate versions and swap these at need, in practice it is complicated as one has to synchronize these versions on the number of already executed stages. Alternatively, we advocate in this paper that the same evolutionary change can be realised by changing the semantics of the loop construct in such a way that it implements the spawning of the stages on the different cores without any further cluttering of the original code. This change would affect a single language construct and is less error-prone when passing from one version to the next, as the change is confined to a single element.

Little to no support for change through language evolution exists in literature in particular if it should be deployed during the application execution. This is mainly due to the fact that language implementations are mostly monolithic and do not support this fine-grained kind of language adaptation but also because an indiscriminate change to a language construct implementation would affect any use of such a construct and this is not always desirable (e.g., when parallelising loops in an application that uses both loops with dependent and independent stages). Therefore, to be effective, such an approach should be supported through an easily evolvable language implementation. Whereas the classic language implementations tend to be monolithic and hard to adapt, in the last decade, several modular development frameworks, such as StrategoXT/Spoofax [11],

---

<sup>1</sup>In this paper, with non-functional we denote a functionality that could be removed without compromising the achievement of the main goal of the application but that represents a sort of general help to get the goal, such as support for concurrency, sustainability and logging.

<sup>2</sup>A loop with independent stages is a loop whose stages do not rely on the results calculated in the previous stage.

Lisa [12] and Neverlang [13] have been developed. These provide a natural environment for the proposed software evolution via language adaptation based approach.

In this paper we propose *to move the evolution from the application level to the programming language level* to support the dynamic addition of non-functional features which originally were not designed into the application. The basic idea is centred around the facts that:

- any application feature is implemented by a subset of the language constructs used to implement that application;
- application feature behaviour is governed by the semantics of the language constructs used in its implementation; and
- the non-functional features introduce additional properties to the existing features (i.e., not stand alone behavior).

So, changing the semantics of the language constructs used in the implementation of the given feature is directly equivalent to changing the feature itself and provides a simple mechanism to extend it with non-functional features. The obvious benefit of our proposed approach is that the application source code will be left untouched. Moreover, the same evolutionary change can be applied to several code portions through a single implementation, potentially reducing code repetition, scope for error introduction, and limiting application code complexity. Furthermore, in a recent work Chitchyan *et al.* [14] introduced the concept of *micro-language*, which is very well suited for handling the change propagation scoping challenge. Micro-languages provide a logical overlay over programming languages, supporting grouping of language concepts in accordance with some application features.

This proposal is supported by a novel DSU framework, called  $\mu$ -DSU, that supports software evolution and, in particular, the addition of non-functional features via language evolution and exploits *micro-languages* to confine the effect of the language changes to only the intended application features. The paper provides three major contributions:

- a clear definition of micro-languages and their relationship with traditional language terminology (Sect. 2);
- the definition of an architecture to support dynamic software updating through changes in the language semantics by using micro-languages (Sect. 3);
- the demonstration of the approach feasibility through the implementation of the  $\mu$ -DSU framework (Sect. 4) and showcasing it with two demonstrative studies (Sect. 5).

With  $\mu$ -DSU a micro-language can be defined for each evolution-prone application feature. As this application feature needs to be evolved this is realised through changes to the implementation of its contained language features. The underlying language implementation for  $\mu$ -DSU is such, that, when the application feature must change, its micro-language can change without side effects on other micro-languages. Thus, the language environment allows each language element to be updated in a localised fashion, with all the change effects absorbed into the evolving element.  $\mu$ -DSU is defined on top of the Neverlang [13] framework for modular language implementation.

In the following, Sect. 2 explains what a micro-language is and its relationship with other programming language concepts. Sect. 3 shows a generic DSU architecture that exploits language adaptation and micro-languages to support dynamic software evolution. Sect. 4 introduces  $\mu$ -DSU that represents a proof-of-concepts for the feasibility of proposed DSU architecture. Sect. 5 shows the micro-language approach to dynamic software evolution on two demonstrative studies. Sect. 6 discusses benefits and drawbacks




Figure 1: Relationship among programming (pl), its sub- ( $sl_i$ ) and micro-languages (the striped blots). Dots represent the different language features.

of this approach, while Sect. 7 presents some related work. Finally, Sect. 8 concludes the paper.

## 2. Introducing Micro Languages.

A micro-language is a logical grouping of constituents from its host programming language determined by the programmatic constructs used in a given piece of code (an application feature). Furthermore, a micro-language is a concept orthogonal to those of language features and sub-languages as it can consist of a number of language features from a number of sub-languages (as depicted in Fig. 1). To illustrate this, we will first present the notions of language features and sub-languages and then show how a micro-language relates to them.

### 2.1. Language Features and Sub-Languages

From a programming language viewpoint, the minimal meaningful concept or construct of a language is called a *language feature* [13]. Examples of language features are the “for” statement, inheritance, or method invocation. As detailed in [13], a *language component* provides the language feature implementation.

Any given programming language can be viewed as the composition of several *sub-languages*, where every sub-language contains a subset of language features of the given (so-called *host*) language. For instance, the SQL language hosts such sub-languages as data definition (which includes, for instance, the create, and drop features), data manipulation (including, for example, select and insert features), transaction boundaries (including, for example, commit and rollback features), etc. In [15], Cazzola and Olivares demonstrated that Javascript can be decomposed into 13 sub-languages.

Sub-languages provide language features to support one (well-defined) *programming aspect* (e.g., setting transaction boundaries in SQL, or dealing with exceptions in Javascript). Yet, to do so they often rely on the presence of other language features provided by other

```

function factorial(n) {
  if (n <= 1) return 1;
  else return n*factorial(n-1)
}

```

Listing 1: Factorial in Javascript

sub-languages (such as creation of a table as a transaction in SQL, or object in Javascript). Therefore, a sub-language needs to be composed with other sub-languages in order to have all its host language features usable. Thus, programming languages could be composed of sub-languages. In either case, a sub-language contains (a number of) language feature(s) that supports one (or more) well-defined aspect(s) of programming. Figure 1 depicts the relationship between language features (the red dots) and a possible set of sub-languages ( $sl_i$ ) of a hypothetical programming language (pl); note that each sub-language restricts the programming language limiting the number of included language features.

## 2.2. Application Features and Micro-Languages

A software application can be described in term of its functionality, often called *application features*.<sup>3</sup> From the application viewpoint, it is irrelevant which language features from which sub-languages are used to implement an application feature. It is rather common to have several language features used in the implementation of a specific application feature, and these language features do not have to belong to a single sub-language. Sub-languages support a single programming aspect but many programming aspects could be required for the implementation of a given application feature.

As an example, let us consider the Javascript implementation for the factorial function (shown in Listing 1). This small piece of code uses—among others—a less equal operator, a conditional statement and a function definition (all language features) belonging to boolean expression, control flow and functional sub-languages respectively. It is evident that this implementation of the factorial application feature relies on (at least) three sub-languages (out of 13 reported in [15]), and each involved sub-language also contains many more language features than those used to implement the factorial. For instance, the boolean expression sub-language also contains all the other comparison operators, not used in this example.

We define a *micro-language* as the set of language features involved in the implementation of a specific application feature. A micro-language is (normally) orthogonal to any sub-language provided by the host programming language. Should we wish to define a micro-language for factorial calculation in Javascript, it would have to utilise language features from at least the above-noted sub-languages.

A micro-language is a logical definition aimed to align an application feature with the language features used in its implementation. Different application features are bound to different micro-languages that could overlap. Two micro-languages  $\mu l_1$  and  $\mu l_2$ —associated with the application features  $af_1$  and  $af_2$  respectively—overlap when there

<sup>3</sup>Czarnecki and Eisenecker [16] defines an application feature as «a distinguishable characteristic of a concept (e.g., system, component, and so on) that is relevant to some stakeholder of the concept.»

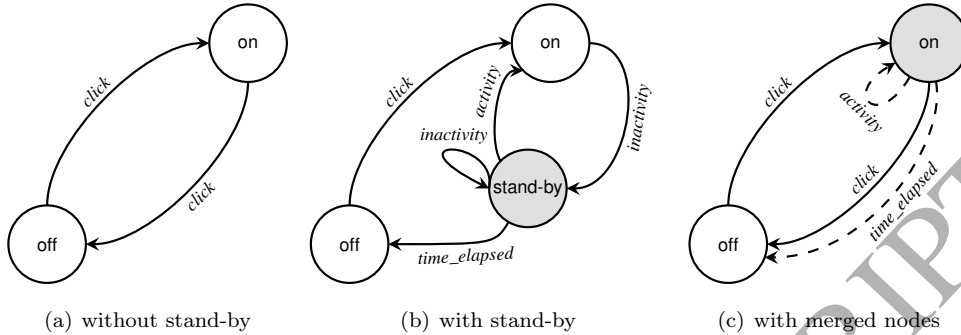


Figure 2: Vacuum Cleaner's default behaviour

exists a language feature  $lf$  that is used to implement both  $af_1$  and  $af_2$  and therefore it belongs to both  $\mu l_1$  and  $\mu l_2$ .

Since a micro-language selectively encapsulates the relevant language features from across relevant sub-languages, a single micro-language will rarely ever be also a usable programming language or be able to construct a full application. Instead, a set of micro-languages will be required to represent the various features of the complete application. Note that the union of such a set of micro-languages represents a full coverage of all language features provided by the host language.

Out of one host programming language, a myriad of different micro-languages could be defined, each crossing sub-language boundaries of the host language. It is important to note that micro-languages do not require any specific interpreters/compiler and they do not have to be composed or interfere with each other but only provide a logical overlay on the used programming language that relates part of it to the application features. In Fig. 1, some micro-languages are highlighted by striped blots.

### 2.3. An Illustrative Decomposition in Micro-Languages

Let us demonstrate the notion of a micro-language with an example for a state machine governing the controller for an automatic vacuum cleaner [14]. The behaviour of any electrical household appliance can be easily modelled with a state-machine: it has some feasible states (e.g., *on*, *off*, etc.) and some transitions that, under given events (e.g., clicks on a button), move the appliance from one feasible state to another. In our case, the default behaviour simply turns the vacuum cleaner on and off when the switch is turned to “on”/ “off” respectively—as depicted in Fig. 2(a).

The code snippet in Listing 2(a) shows the implementation for the default behaviour of the vacuum cleaner (leaving out the code for moving the vacuum cleaner that it is not relevant to the discussion). At least, two application features are implemented: i) the *turning on* of the vacuum cleaner and ii) the *turning off* of the vacuum cleaner. These two application features are bound to two micro-languages denoted by  $\mu_{ton}$  and  $\mu_{toff}$  respectively. Turn on for the vacuum cleaner consists of changing the *off* state into the *on* state when a *click* event occurs. Looking at the language:

- the definition of a particular *state* is governed by the *state initialization* language feature identified by the **states** declaration statement;



<pre> <b>states</b>   on = <b>turn-on</b>   off = <b>turn-off</b> <b>events</b>   click = get_click() <b>transitions</b>   on { click ⇒ off }   off { click ⇒ on } </pre>	<pre> <b>counter</b> t <b>states</b>   on = <b>turn-on</b> ; t ← 0   off = <b>turn-off</b>   stand-by = t ← t + 1 <b>events</b>   click = get_click()   elapsed = t &gt; 10   activity = get_activity()   inactivity = get_time() <b>transitions</b>   on {click ⇒ off ; inactivity ⇒ stand-by}   off {click ⇒ on}   stand-by {     click ⇒ off ; elapsed ⇒ off ;     inactivity ⇒ stand-by ; activity ⇒ on   } </pre>
---	--

(a) default behaviour

(b) extended behaviour

Listing 2: Vacuum Cleaner's Code

- the initialization of the *on* state is governed by the *turning on* language feature identified by the **turn-on** operation;
- the transition from one state to another is governed by the *transition definition* language feature identified by the **transitions** declaration statement.

Similar considerations apply to the *turning off* application feature with the only difference that the state to be defined is *off* and this is governed by the *turning off* language feature. Thus, the language features describing the two micro-languages are:

$$\begin{aligned} \mu_{ton} &= \{ \textit{turning on}, \textit{state initialization}, \textit{transition definition} \} \\ \mu_{toff} &= \{ \textit{turning off}, \textit{state initialization}, \textit{transition definition} \} \end{aligned}$$

From this example we see that:

- micro-languages differ from the language used to implement the behaviour of the vacuum cleaner
- a micro-language could be unusable, e.g., neither  $\mu_{ton}$  nor  $\mu_{toff}$  include the language features necessary to define the available or specific events without which it is impossible to define a working state machine.
- micro-languages can overlap.

Thus, to re-cap: micro-languages overlay a logical structure on the *host language*; each of them defines and governs the behaviour of a given application feature and straddles several sub-languages from those provided by the host language.

### 3. Evolution via Micro-Languages

An application feature is bound to the set of language features needed to implement it that are grouped by a micro-language. Thus, any change that should affect an application feature can be governed and deployed through changes to the respective micro-language.

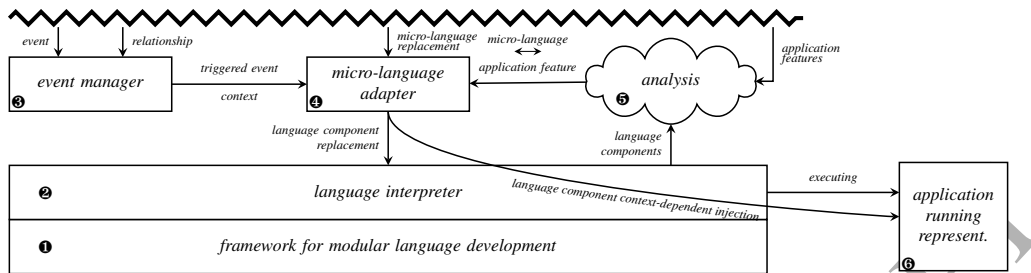


Figure 3: Architecture for Micro-Language Based Adaptation of a Running Application

### 3.1. Driving the Adaptation

Figure 3 graphically overviews the architecture of a DSU framework that would adapt running applications through the adaptation of micro-languages. A modular language development framework (marked with ① in Fig. 3) lays the foundation for this architecture and furnishes the necessary mechanism to deal with the micro-language adaptation. The modular language interpreter of the chosen language (marked with ② in Fig. 3)—e.g., that of the state machine language used in Sect. 2.3—runs the application that should be evolved (marked with ③ in Fig. 3)—e.g., the vacuum cleaner control application. Each micro-language will correspond to one application feature. Currently the application features are provided by the application developer and it is represented by a set of coordinates (code lines and element description) to the application source code. A way to automatically elicit from the application analysis both the application features and the language features used in its implementation—and therefore the micro-language associated to the application feature—is currently under investigation. This activity is represented by the component marked with ⑤ in Fig. 3.

The evolution of an application feature is implemented through the evolution of the corresponding micro-language. Each change to a language feature in the micro-language affects the application feature behaviour. The adaptation process is event-driven; the *event manager* component (marked with ③ in Fig. 3) waits for the (application specific and user-defined) *update events* whose occurrence will trigger the adaptation process. The event manager component deals both with the update events it is to monitor and with the associations between these events and the application features (and hence also the micro-languages) interested in the occurring events. The relationship between occurring events and changes to the application features along with micro-languages, provides the system with the *context* information necessary to limit the effect of the micro-language adaptation. For instance, in the vacuum cleaner example, the inactivity of the running appliance is the event that triggers the adaptation, thereby affecting the *turning on* application feature. The *micro-language adapter* component (marked with ④ in Fig. 3) is in charge of the adaptation process.

The change at language level, if not limited, will affect the behaviour of the whole application. This can be a useful trait in cases where the required change is general: it has the obvious benefit of a limited adaptation instead of looking for, and individually changing each part of the relevant application features. On the other hand, often an adaptation is desired for a specific, localised, case and such a change should affect only a few application

features. Thus, the language adaptation-based framework has to support both system wide and selectively targeted exchange of language features; both types of changes should be done without forcing a regeneration of the interpreter or a mandatory stopping and restarting of the running application. Let us demonstrate these two adaptation scopes using the vacuum cleaner example presented in Sect. 2.3. The two micro-languages  $\mu_{ton}$  and  $\mu_{toff}$  share the *state initialization* language feature implemented by the `State` slice. A system wide change at the interpreter-level to the *state initialization* language feature would affect both micro-languages and the associated application features. Instead, a selective exchange of language features would permit a fine-grain control of the adaptation process. It would permit to subordinate how a language feature (*state initialization* in our case) is interpreted according to the belonging to a specific micro-language. The association of an application feature to its AST permits to directly inject such a semantic action only in the nodes relevant to the  $\mu_{ton}$  micro-language and to avoid any interference with the same type of nodes that are associated to the  $\mu_{toff}$  micro-language. This is possible because the application feature to adapt is represented by a specific AST portion; the association with a micro-language permits to identify such an AST portion and to limit the effect of the adaptation only to that portion, granting both a fine-grain control and a context-dependent language adaptation. Therefore, given the context, the adapter will not modify the interpreter, but instead will operate directly on the AST by modifying the semantic actions that should be executed when specific nodes of the tree are visited. Details about how this should work are explained in Sect. 4.

### 3.2. Evolution of Micro-Language Semantics

Let us consider the scenario where the requirements are evolved in the vacuum cleaner example presented in Sect. 2.3. The new version of the system requirements pay more attention to reduction of energy waste. Now the requirement for switching the system on/off is updated to stating that (for safety and to reduce energy waste) the controller should switch the vacuum cleaner off when it is not in productive use (e.g., when it is accidentally switched on by a toddler). By definition, when in productive use the vacuum cleaner should be in motion; if it is not in motion for 10 seconds, it is considered not in use.

Supporting the behaviour of “turning the vacuum cleaner off when inactive for a given period” at the code-level would require (Fig. 2(b)) addition of new:

- state `stand-by` used when the appliance is on but inactive,
- event `time-elapsed` which is triggered when the specified number of seconds have elapsed without any activity,
- transition from the `stand-by` state to `off` state when the `time-elapsed` event occurs, and
- events handling activity/inactivity of the vacuum cleaner.

Although such code-level adaptation is quite simple and general, it is also very invasive and thus not advisable if it is to be applied to a large set of various types of electronic appliances. The adapted code is shown in Listing 2(b).

When using the micro-languages, instead of modifying the controller program with condition checks for motion monitoring (as per Listing 2(b)), we can change the semantics of some of the language constructs. We call such adaptation *seamless*<sup>4</sup>. In this example

<sup>4</sup>This is in contrast to an *explicit adaptation* that introduces new syntax and forces an application

the meaning of being in the *on* state can be changed by re-defining the semantics of the *turning on* language feature. In short, the **turn-on** operation will not only turn on the appliance but also store the timestamp of when it is turned on. Seamlessly adapted controller behaviour is represented by the state machine in Fig. 2(c) where two (internal) events: `activity` and `time_elapsed` are added, as well as transitions (dashed in the figure to note the fact that they are a sort of a side effect of the changes to the language semantics) from the state `on` to the state `on` (guarded by `activity`) and to state `off` (guarded by `time_elapsed`). The former event will reset the stored time as a consequence of some activity, while the latter will check the current time against the stored time to detect if it is time to turn the appliance off due to prolonged inactivity.

Such a *seamless* adaptation is enabled by changing the implementation of the  $\mu_{ton}$  micro-language and will leave the controller code unchanged. Therefore, any program written in this language will remain unaltered, yet will incorporate the new behaviour. Moreover, a micro-language provides a context in which the language feature is used; and this context can be used to limit the effect of the change to a specific occurrence of the language feature. Furthermore, if the host programming language is based on a framework for modular language development [12, 13, 17], adapting a micro-language is only a matter of plugging/unplugging a few language features and in some cases this can be done without regenerating the interpreter.

As shown, language adaptation via micro-languages can ease the application's adaptation. Here the adaptation is moved from the application to the language implementation.

#### 4. $\mu$ -DSU: Implementation Details

As a proof-of-concept, the architecture in Fig. 3 has been implemented into the  $\mu$ -DSU framework over Neverlang [13].

To provide an usable micro-languages-based approach to software evolution, we must provide an operational environment for it. Modular language development frameworks [11–13, 18] emphasise the separation of language features as pluggable and composable units (slices in Neverlang) and they represent a perfect fit when the composition can occur during the application interpretation such as in Neverlang [19, 20].

Note that the micro-language concept does not have a direct match in any modular language development framework but it can be defined within the framework concepts. In the Neverlang parlance, a micro-language is a logical cluster of some of the slices used in the interpreter for the host language.

##### 4.1. Neverlang in a Nutshell

The *Neverlang* [13, 21–23] framework is built around the language feature concept. Language components, called *slices*, embodying the language features are developed as separate units that can be compiled and tested independently, enabling developers to share and reuse the same units across different language implementations. Here the development base unit is the `module` (Listing 3). A module may contain a `syntax` definition and/or a `semantic` role. A role defines actions that should be executed when some syntax is recognized, as prescribed by the *syntax-directed translation* technique [24]. Syntax

---

change. This is not further discussed for clarity.

```

3  module StateModule {
4      reference syntax {
5          State StateName "=" Expr
6      }
7      role(execution) {
8          0 { newstate($1.state, $2.action); }.
9      }
10 }
11
12 slice State {
13     concrete syntax from StateModule
14     module StateModule with role execution
15 }
16
17 language HooverLang {
18     slices Program StateDecl EventDecl TransDecl StateLst State StateName
19         EventList Event EventName Expr BExpr TransList Transition Support
20     roles syntax < execution
21 }

```

Listing 3: Slice the syntax and semantics for the state concept.

definitions and semantic roles are tied together using `slices`. Let us see the Neverlang realisation of the `State` module for the vacuum cleaner example shown in Listing 3. Here the module `StateModule` declares a reference syntax for the state concept (lines 2-4) and actions are attached to the nonterminals on the right of the production (line 7). Semantic actions are attached to nonterminals by referring to their position in the grammar or through a label: numbering starts with 0 from the top left to the bottom right<sup>5</sup>, so the first `State` on line 3 is referred to as 0, `StateName` as 1, and the `Expr` is referred as 2. The slice `State` declares in line 12 that we will be using *this* syntax (which is the concrete syntax) in our language, with those particular semantics (line 13). Finally, the language descriptor indicates (lines 17-18) which slices are to be composed together to generate the language interpreter. Composition in Neverlang is, therefore, twofold: (1) between modules, which yields slices, and (2) between slices, which yields a language implementation. The composition result is independent of the order of specified slices. The grammars are merged to generate the complete language parser. Semantic actions are performed with respect to the parse tree of the input program; roles are executed in the order specified in the `roles` clause of the `language` descriptor. Please see [13] for further details.

#### 4.2. $\mu$ -DSU over Neverlang

The Neverlang [13] framework permits changes to the interpreted language without regenerating its parser [25]. Changes can be introduced during the interpretation of the application both at system-wide level and localised on the application *parse tree* (PT) [20, 26]. These Neverlang's characteristics allow  $\mu$ -DSU to support dynamic language adaptation in two ways: i) by replacing language components and ii) by directly modifying how the language feature is interpreted. The former simply consists of replacing a

<sup>5</sup>Neverlang also permits to label a production and refer nonterminals via an offset from such a label, e.g., `$STATE[0]` is the head of the `STATE` production.

Context Definition
<b>[endemic] slice</b> «id <sub>1</sub> » [, «id <sub>2</sub> », ...] : «slc» ; <i>To bind the (endemic) slice «slc» to a name «id<sub>1</sub>»; if multiple names are provided they are all aliases for the same (endemic) slice.</i>
<b>production</b> «id <sub>1</sub> » [, «id <sub>2</sub> », ...] : «rule» <b>from module</b> «mod» ; <i>To bind a rule «rule» from a module «mod» to a name «id<sub>1</sub>»; if multiple names are provided they all refer to the same rule.</i>
<b>nt</b> «id <sub>1</sub> » [, «id <sub>2</sub> », ...] : «rule» <b>from module</b> «mod» ; <i>To unpack into «id<sub>1</sub>», «id<sub>2</sub>», ..., «id<sub>n</sub>» the first n nonterminals in «rule» from the module «mod».</i>
<b>action</b> «id» : «nonterminal» <b>from module</b> «mod» <b>role</b> «name» ; <i>To bind the action associated to the «nonterminal» from the module «mod» to the name «id».</i>
Matching Expressions
«id»[[«cond <sub>1</sub> (attr <sub>1</sub> )» [, «cond <sub>2</sub> (attr <sub>2</sub> )», ...]]] <i>Matches the PT node identified by «id» when its (specified) attributes («attr<sub>1</sub>») verify some conditions («cond<sub>1</sub>()»). The conditions are predicates that compares the current value of the attribute against a constant; if no condition is given the node is matched by name.</i>
«id <sub>1</sub> »[[«cond(attr)»]] < «id <sub>2</sub> »[[«cond(attr)»]] [   «id» ] <i>Matches a path on the PT where the node «id<sub>1</sub>» is the parent of the node «id<sub>2</sub>». As before conditions on the node attributes can be given and the filter operator ( ) permits to bind the result to the desired node otherwise it will be the parent.</i>
«id <sub>1</sub> »[[«cond(attr)»]] << «id <sub>2</sub> »[[«cond(attr)»]] [   «id» ] <i>Matches a path on the PT where the node «id<sub>2</sub>» can be reached from the node «id<sub>1</sub>». As before, it is possible to express conditions on the node attributes and a filter operator.</i>
Manipulation Operations
<b>add action</b> «id <sub>1</sub> » [to «id <sub>2</sub> »] <b>in role</b> «name» ; <i>To add the action «id<sub>1</sub>» to the node «id<sub>2</sub>» in the PT for the role «name»; if the target node is omitted the one matched is used.</i>
<b>remove action</b> «id <sub>1</sub> » [from «id <sub>2</sub> »] <b>in role</b> «name» ; <i>To remove the action «id<sub>1</sub>» from the node «id<sub>2</sub>» in the PT for the role «name»; if the target node is omitted the one matched is used.</i>
<b>set specialized action</b> «id <sub>1</sub> » [to «id <sub>2</sub> »] <b>in role</b> «name» ; <i>To set the specialized action for the nonterminal «id<sub>2</sub>» to «id<sub>1</sub>» in the role «name»; if the target node is omitted the one matched is used.</i>
System-Wide Manipulation Operations
<b>replace slice</b> «id <sub>1</sub> » <b>with</b> «id <sub>2</sub> » ; <i>To replace the slice «id<sub>1</sub>» with slice «id<sub>2</sub>».</i>
<b>redo</b> [from «node»] [in role «name»] ; <i>To restart the visit of the role «name» from a given node «node» as expressed in the context; the visit restarts from the root of the current role when «node» and «name» are omitted respectively.</i>

Table 1: Summary of the  $\mu$ DA DSL.

language component (a slice in Neverlang parlance) affecting every use of such a language feature. The latter permits to programmatically modify—according to the micro-language definitions—a language feature locally to a single use of such a feature by altering its behaviour according to the context provided by the PT [20].

Adaptation rules for the micro-languages are expressed through a dedicated *micro-dynamic adaptation* domain specific language (DSL):  $\mu$ DA.  $\mu$ -DSU implements a *micro-language adapter* (component ④ in Fig. 3), which executes rules written in  $\mu$ DA, and interacts with the interpreter (component ② in Fig. 3) written in Neverlang (component ① in Fig. 3) during the application execution. The event manager runs a script that registers the set of possible events and the adaptation scripts that need to be activated in case of a given event occurrence. The event register script has to be maintained by the application developer. The *event manager* component (component ③ in Fig. 3) notifies the adapter about the events for which specific micro-languages have registered interest. The adapter

triggers the respective scripts to deploy the required adaptations for each event per each micro-language.

A generic  $\mu$ DA script has two sections. The first section defines the context for the script, where the PT nodes are described in term of their original productions or in terms of the nonterminal these nodes define; these definitions are bound to names so that they can be referenced through these names elsewhere in the script (the set of constructs used in the context definition is shown in Tab. 1). Each context mirrors a micro-language associated to the given application feature and all the elements of the language features of that micro-language are accessible in the script's context (e.g., see Listing 8(a) lines 1-2). In the context section the actions that should be injected are also defined (e.g., see Listing 8(a) line 4). The code to be injected is written in Neverlang slices which serve as containers for new, replacement adaptation actions. In the case of localised manipulations, the second section declares a set of adaptation clauses to be applied to a specific portion of the PT. A sample structure of such clauses is shown below:

```
when «matching expression» occurs {«manipulation operations»}
```

The matching expression resembles an AspectJ's pointcut [27] that match against paths of the PT's application. The matching expressions—described in Table 1—permit to identify a node by its kind (as defined in the context section), by the values of its attributes and by node's relative position in the PT (e.g., if it is a parent of another node). Upon the request from the event manager to the adapter (see Fig.3), the expression is matched against the whole PT. According to the Neverlang's architecture for dynamic adaptation [20], an agent—whose body is composed of the corresponding manipulation operations—is attached to the matched node and its code will be executed immediately before the next visit of the node that will occur during the interpretation and after that removed. The supported manipulation operations permit to add/remove an action to/from a node with respect to a specific role<sup>6</sup>. In the case of a system-wide adaptation, the manipulation is independent of the PT but it directly affects the slices definitions. Therefore no match against the PT is needed and the manipulation operations are simply introduced by the keyword **system-wide**. At the moment, only the slice replacement and the restarting of the PT visit are supported. Table 1 summarises the whole set of operations belonging to  $\mu$ DA; examples of its use are showed in Sect. 5.

The  $\mu$ -DSU implementation of the *event manager* is quite simple and its behaviour is governed by bash scripts. The manager is a process that waits for external events or from the Neverlang virtual machine. Each event blocks the interpretation of the application. If the event does not require an adaptation the event manager immediately acknowledges the event and the interpretation resumes. Where adaptation is required, the event and its matching contexts (from the registry script) are passed to the micro-language adapter. When the adapter completes the execution of the adaptation scripts, it notifies the Neverlang virtual machine and the interpretation is resumed from the point of the previous suspension, but using the adapted interpreter.

<sup>6</sup>Neverlang supports several user-defined roles that correspond to consecutive visits of the PT and at each visit, a different action can be executed.

Language feature	Description
<code>print Expr</code>	prints Expr to screen
<code>set font size Expr</code>	sets the font size to Expr points
<code>set font color Expr</code>	sets the font colour to Expr

Table 2: The language features added to Neverlang.JS.

## 5. Demonstrative Studies

In this section, we showcase feasibility of  $\mu$ -DSU through two demonstrative studies.<sup>7</sup> In the first study, we show how an HTML viewer application seamlessly adapts its configuration to the accessibility needs of the current user. This study showcases a system-wide change, where the change of the evolving micro-language does not create an interference with other micro-languages. The introduced non-functional feature is the accessibility support. In the second study, we speed up the Mandelbrot set drawing by parallelising its computational engine. This study shows a case of local adaptation, where the evolving language features in one micro-language is also part of another micro-language that should stay untouched, so the change must be localised. The introduced non-functional feature is optimisation. Together these studies demonstrate the main characteristics of our approach.

### 5.1. System-Wide Adaptation

Consider the benefits of having an HTML viewer that adapts page visualisation to the conditions of the user’s eyesight. For non-impaired users, the viewer would render the page as specified by the underlying HTML code. For users suffering from colour-blindness green and red colours would be substituted; for long-sighted users, a larger font would be used; for blind users, the displayed text would also be read aloud. In several occasions this kind of application would need dynamic adaptations, e.g., digital information booth at a convention where several people with different eyesight needs will continuously alternate or a mobile device that can be passed to a friend with different needs.

As a proof-of-concept, a simple HTML viewer has been developed and  $\mu$ -DSU is used to dynamically adapt such a viewer to the current user’s eyesight. This example demonstrates how micro-languages ease *a posteriori* development and deployment of application variants according to dynamically changed needs either neglected at design time—e.g., due to higher expenses—or newly identified. The viewer has been written in Neverlang.JS [15]: a Neverlang implementation of the Javascript interpreter. Neverlang.JS has been extended with three language features (summarised in Tab. 2) that help to illustrate and simplify the definition of micro-languages.

Central to the HTML viewer is the application feature of *show text on screen* ( $AF_{show}$ ) responsible for the HTML text visualization.  $AF_{show}$  is implemented by the function `view()` (Listing 4) and uses the language features summarised in Tab. 2. Thus, these three language features effectively capture a micro-language for the  $AF_{show}$  application feature:

<sup>7</sup>The demonstration of both studies is shown in the movie:  
[http://cazzola.di.unimi.it/ \$\mu\$ -dsu/ \$\mu\$ -dsu-demo.mp4](http://cazzola.di.unimi.it/<math>\mu</math>-dsu/<math>\mu</math>-dsu-demo.mp4).



$$\mu_{show} = \{ \text{set font color, set font size, print} \}$$

The original application does not consider any accessibility issue. The initial  $AF_{show}$  application feature simply parses the HTML element which extracts the colour, size and text from it; then, the colour and size font properties are set according to the parsed values and finally the text is printed. The slice implementing the original behaviour of the print language feature of the  $\mu_{show}$  micro-language is shown in Listing 5(a). As it becomes necessary to support visually impaired users we change the way that  $AF_{show}$  visualises the HTML text by changing the semantics for the language element of the  $\mu_{show}$  micro-language. Listing 5(b) shows the implementation for the print language feature of  $\mu_{show}$  intended for the long-sighted users. The semantic action, among other things, extracts the current size setting which is then multiplied by 3 to increase the size of the visualised text. Similarly, Listing 5(c) shows the implementation of the print for blind users. Besides printing the text, this slice also reads it aloud using a specific library to read text. Note that the module does not define any syntax. In fact, the `slice` construct permits to reuse the original `print` syntax defined in Listing 5(a) and to link it to the semantic action for blind users. To save space and since it will not add anything to the discussion, the slice for colour-blind users is not shown.

To support any of the visually impaired users, we simply integrate the relevant changes to the given micro-language in the interpreter by replacing the appropriate modules on need. The original application code never changes, i.e., the application is unaware of the evolving requirements. Moreover, the change is confined in one point—the implementation of the `print` feature—that eases its management. Figure 4 shows HTML page screenshots visualised with the slices for normal, long-sighted, and colour-blind users (the case of blindness is not shown for obvious reasons).

Let us detail how the process explained in Sect. 3.1 and summarised in Fig. 3 is applied to this demonstrative study. To allow for runtime adaptation based on user’s (dis)ability, (separately from the viewer application) user profiles are associated with their log-in accounts. Let us assume that the user profile change is identified according to the fingerprint/login id provided by the new user when starting to look at the viewer that is already in use. Note, the *profile* changes at system level not at the application level; the application still behaves as usual. The profile change is the event that triggers the adaptation of the application. When the user profile changes, the event manager notifies the adapter component about the new profile and the concerned context which in our case is the  $AF_{show}$  application feature. Finally, according to the new profile, the slices associated with the language features in the  $\mu_{show}$  micro-language are replaced with the version corresponding to the visual impairment of the current user. In this case, the change will affect all uses of the changed language features since the same kind of accessibility support is needed for the whole application. Notice that when, in response to a profile change event, we switch slices in the interpreter or we inject semantic actions

```
function view(element) {
  var parsed = parse(element); set font color parsed.color;
  set font size parsed.size; print parsed.text;
}
```

Listing 4: JavaScript snippet of the view application feature.

```

module sustainability.HealthyPrint {
  reference syntax {
    Print: Statement ← "print" Expression;
  }
  role (evaluation) {
    Print: @(
      JSType t = (JSType)$Print[1].value;
      int size = view.getCurrentSize();
      view.print(t.stringValue());
    ).
  }
}
slice Print {
  concrete syntax from sustainability.HealthyPrint
  module sustainability.HealthyPrint
  with role evaluation
}

```

```

module sustainability.HyperopicPrint {
  reference syntax from sustainability.HealthyPrint
  role (evaluation) {
    Print: @(
      JSType t = (JSType)$Print[1].value;
      int size = view.getCurrentSize();
      view.print(t.stringValue(), size+3);
    ).
  }
}
slice HyperopicPrint {
  concrete syntax from sustainability.HealthyPrint
  module sustainability.HyperopicPrint
  with role evaluation
}

```

```

module sustainability.BlindPrint {
  reference syntax from sustainability.HealthyPrint
  role (evaluation) {
    Print: @(
      JSType t = (JSType)$Print[1].value;
      view.print(t.stringValue());
      view.speak(t.stringValue());
    ).
  }
}
slice BlindPrint {
  concrete syntax from sustainability.HealthyPrint
  module sustainability.BlindPrint
  with role evaluation
}

```

(a) default print slice

(b) print slice for long-sighted users

(c) print slice for blind users

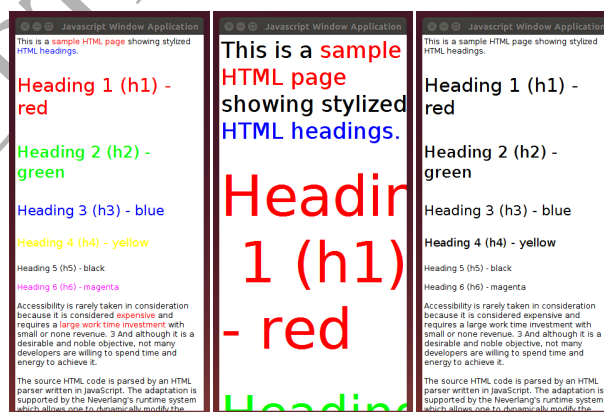
Listing 5: Three versions of the implementation for the print language feature.

into the syntax tree, the original application code never changes, i.e., the application remains unaware of the supported adaptation.

Listing 6(b) shows the  $\mu$ DA script for swapping the current slice for the print language feature with the one needed when the user's profile change from `healthy` to `blind`. Notice that, the change is driven by the context definition (the new and old names) that is provided by the association between the  $\mu_{show}$  micro-language and the  $AF_{show}$  application feature. The micro-language provides information about which language features (slice in the Neverlang parlance) should be affected by the change whereas the application feature describes where such language features are used limiting, de facto, the effect of the change. Listing 6(a) shows the script that drives the adaptation by executing the  $\mu$ DA script in Listing 6(b) (line 5) when the event occurs (i.e., when the user profile changes).

## 5.2. Localised Adaptation

The second demonstrative study considers the calculation and pictorial drawing of the Mandelbrot set, aiming to speed up the set calculation by parallelising the computation cycles [9, 28, 29]. As discussed below, we define micro-languages for the application features and the parallelising change is supported through these micro-language change.



(a) normal

(b) long-sighted

(c) colour-blind

Figure 4: A sample page displayed according to 3 of the possible profiles.

<pre> current="healthy" while true; do 3   profile=getprofile()    if ["\$profile" != "\$current"]; then 6     µda "\$current→\$profile.µda"      current="\$profile"    fi    sleep 1 9 done </pre>	<pre> slice old: sustainability.HealthyPrint; slice new: sustainability.BlindPrint; 3 system-wide {    replace slice old with new; 6   redo role evaluation; 9 } </pre>
(a) event manager script	(b) healthy→blind.µda

Listing 6: Scripts for the HTML viewer management

```

var MAX_ITER = 50; var ZOOM = 450; var HEIGHT = 300; var WIDTH = 300;
var I = create2DArray(WIDTH,HEIGHT);
3 var zx=0; var zy=0; var cX=0; var cY=0; var tmp=0; var iter=0;
for(var y = 0; y<HEIGHT; ++y) {
   for(var x = 0; x<WIDTH; ++x) {
6     zx = 0; zy = 0; cX = (x - 400) / ZOOM;
     cY = (y - 300) / ZOOM; iter = MAX_ITER;
     while (((zx * zx) + (zy * zy)) < 4) && (iter > 0) {
9       var tmp = (zx * zx) - (zy * zy) + cX;
         zy = (2.0 * (zx * zy)) + cY; zx = tmp; iter = iter - 1;
     };
12    I[x][y] = iter | (iter << 8);
};
};

```

Listing 7: Javascript for the Mandelbrot set calculation.

Listing 7 shows a traditional Javascript implementation of the Mandelbrot set calculation algorithm. Several application features can be identified in this code, such as:  $AF_{calc}$  that calculates the *values* of a point (covering the `while` statement at line 8 with its body),  $AF_{cols}$  that cycles on the columns of the final picture (the nested `for` at line 5 with its body but excluding the calculation of the point) and  $AF_{rows}$  that cycles on the rows of the final picture (the `for` statement at line 4).

Micro-languages bound to  $AF_{calc}$ ,  $AF_{cols}$ , and  $AF_{rows}$  (with the language features denoted by their syntactic symbols) are:

$$\mu_{calc} = \{ while, =, -, +, *, >, \&\& \}$$

$$\mu_{cols} = \{ for, =, -, /, <, ++ \}$$

$$\mu_{rows} = \{ for, =, <, ++ \}$$

As can be noticed above, these three micro-languages are largely overlapped, so a change to one of the language features (e.g., to `for` or `=`) will also affect the other application features if performed system-wide.

The two nested `for` loops in Listing 7 clearly have independent stages (the values of each point are calculated fully independently) and so could be both parallelised. This could be easily done by replacing the slice for the `for` statement with one with a parallel implementation and immediately all the occurrences of the `for` would be parallelised (similar to the system-wide adaptation of the `print` language feature in the previous study




Figure 5: Selective for adaptation

(Sect. 5.1)). But a more interesting case is to see if the change could be localised to *only one* occurrence of the `for` statement—let’s say, the nested occurrence which is used by the  $AF_{cols}$ —application feature leaving the other occurrence unchanged.

Figure 5(a) visualises a portion of the PT for the Mandelbrot set calculation showed in Listing 7; the nodes representing `for` statements are purple coloured. This PT is rooted at the node for the first `for` statement (line 4). This `for` statement has four parameters: initialization, condition, increment, and body, all of which are referred to by the node’s children. For sake of simplicity, in Figure 5(a) we draw only two children of this node: the node on the left represents the first three parameters and the tree on the right represents the body of the `for` statement. The body of the `for` statement is a sequence of statements so its first node represent the sequence operator (the `;`)<sup>8</sup>. The second `for` statement (line 5) is represented by the left sub-tree of the sequence operator. The rest of the tree is irrelevant for this discussion and not further considered.

Looking at the PT gives an at-a-glance view of the relations between different statements during the interpretation. It is evident that there is a path in the tree from the node of the first `for` to the second one. Similarly, the definition of the application feature for  $AF_{cols}$  and therefore the belonging to the  $\mu_{cols}$  micro-language provide the same information. Note that an application feature is demarcated as a set of code lines/statements. In Fig. 5(a), the  $\mu_{cols}$  micro-language is represented with a striped bloat. The effect of substituting any occurrence of the `for` language feature with a parallel `for` is illustrated in Fig. 5(b). In this case, all nodes representing `for` occurrences (the purple ones) are replaced by nodes representing parallel `for` occurrences (the light green ones). Instead to get the desired evolution (Fig. 5(c)) with only the second occurrence of the `for` language feature replaced we have to exploit the context information provided by the  $\mu_{cols}$  micro-language as done in the  $\mu DA$  scripts reported in Listing 8 where the `«` operator (line 6) is used to capture the second `for` occurrence.

The whole experiment binds the speed up/slow down of the Mandelbrot set calculation to the presence/absence of a stable power source. When the laptop runs on the battery it uses only one core to save the battery and the sequential implementation for the `for` statement is used. When the laptop is plugged into the mains, more cores can be used and the sequential `for` is replaced with the version that executes the loop stages in parallel

<sup>8</sup>This is true also when the body is a single statement because Neverlang.JS, for sake of simplicity, does not have the simplification rule for this case.

on more cores. The event manager monitors the battery and toggles between the two scenarios (script in Listing 9). Two  $\mu$ DA scripts are used (Listings 8(a) and 8(b)) to deal with the toggle of the language features. Both rely on the context information that comes from the micro-language  $\mu_{cols}$ . The nonterminals (the head in particular) of the production for the `for` statement are retrieved and used to search the PT for its occurrences. Similarly, the action to inject updated code is retrieved from a given slice and used where necessary, in accordance with the path matching operators. The only real difference between the two scripts is the action to be injected, toggling between the sequential and parallelised scenarios.

To informally evaluate the effectiveness of our run-time adaptation solution, we get the program to calculate a Mandelbrot set of  $300 \times 300$  points, that is, the external `for` loops 300 times. We run the application in 3 batches of 100 runs each, starting with an unplugged (i.e., single core) state. As the application is running, we plug in the laptop, thus triggering adaptation to multi-core calculation. In the first batch of runs, the laptop was plugged in between the 0th and the 75th lap of the external loop and we obtained a speed up of 44.32% on average. In the second batch of runs, the laptop was plugged in between the 76th and the 150th lap of the external loop and the achieved speed up is 31.32% on average. In the last batch of runs, the laptop was plugged in between the 151st and the 225th lap of the external loop and a speed up of 20.65% on average was achieved. Plugging the laptop after the 225th lap has been considered insignificant from the point of view of the speed up. As expected, the sooner the parallel implementation of the `for` statement was used in the running application, the larger was the achieved speed up.

## 6. Discussion

Micro-languages are enablers for dynamic software evolution. Here we discuss the strengths and weaknesses of this approach and we provide an initial evaluation of its benefits.

```

3 nt for1,_,_,_ : For from module neverlangJS.ForLoop;
  nt for2,_,_,_ : For from module neverlangJS.ForLoop;
  action parforAct : ParFor from module neverlangJS.ParForLoop role evaluation;
6 when for1 << for2 | for2 occurs {
  set specialized action parforAct to for2 in role evaluation;
}

```

(a) `for`→`parfor`. $\mu$ da

```

3 nt for1,_,_,_ : For from module neverlangJS.ForLoop;
  nt for2,_,_,_ : For from module neverlangJS.ForLoop;
  action forAct : ParFor from module neverlangJS.ForLoop role evaluation;
6 when for1 << for2 | for2 occurs {
  set specialized action forAct to for2 in role evaluation;
}

```

(b) `parfor`→`for`. $\mu$ da

Listing 8:  $\mu$ DA's scripts to toggle `for` and `parfor`.

```

[[ $(acpi -a) == *"off"* ]]; plugged=$?
while true; do
  [[ $(acpi -a) == *"off"* ]]; now=$?
  if [ "$plugged" != "$now" ]; then
    if [ "$now" == true ]; then
      # state changed from plugged to unplugged
      echo "Power saving mode enabled";
      µda parfor→for.µda
    else
      # state changed from unplugged to plugged
      echo "Performance mode enabled"
      µda for→parfor.µda
    fi
    plugged="$now"
  fi
  sleep 1
done

```

Listing 9: Event manager script to deal with power source.

### 6.1. Modularity through Micro-Languages.

By definition, a micro-language is conceptually aligned with an application feature and so directly encapsulates the feature concern. Hence, the adaptation of an application feature implies a very localised update to the related micro-language and, to (all or some) occurrences of the corresponding language features. In many cases, the update would affect only a few semantic actions. Furthermore, the modular approach that associates an application feature with a micro-language both limits the effect of the change precisely to the relevant application feature, and confines the required verification/validation for the new code to that single feature. This allows for better scoped and effective verification/validation, with no need to consider the scope of the whole application.

It is also notable that this approach also separates the syntactic definition of control flow of the program from the semantics of the executed computation. For instance, in the above presented examples the same given `print` syntactic statement results in different visualisations of the HTML page, and the given `for` structure of the Mandelbrot set algorithm results in different computation sequences. Thus, a new modularity construct for concern separation at the application feature level is provided.

In short, micro-languages foster modularity as: i) they are built from language features that are modularly implemented; ii) each language feature provides a clear composition interface that allows limiting the ripple effects [30] of changes; and iii) a language feature can be implemented and validated separately of the application that will use it.

### 6.2. Composability of Micro-Languages.

While the modularity support of the micro-languages is clear and beneficial, it cannot work without adequate composition support. In this approach composition is carried out through micro-language definition and interpretation at event occurrences. Here: i) each micro-language relates to the events that are relevant to it; ii) each micro-language defines what will be executed by re-defining the meaning of the language features available in the given language. The composition is the most challenging part of the implementation and it strongly depends on good tool support for correctness checks, usability, and general

	application-level		micro-language	
	adaptation	variants	language feat.	applicat. feat.
HTML Viewer	26	215	57	0
Battery	12	64	48	0

Table 3: Number of line of code changed for adaptation

acceptance. Being a rather novel proposal, at present, this approach has little to offer in terms of development tool support. This can be particularly difficult to handle when the adaptation effects need to be limited by the context provided by the micro-languages. In the absence of good tool support, a misalignment between the application features and the micro-languages could occur and it could have devastating effects with undesired application adaptations due to an unmarked overlapping of micro-languages. Furthermore, a change at language level could introduce new potential vulnerabilities and security risks—a topic that has yet to be investigated.

### 6.3. On the Ease and Cost of Adapting Application Features.

In order to validate the potential utility of our approach, we must discuss if the effort needed for dynamically adapting a given application through micro-languages is less than that needed for adapting that same application code without the help of micro-languages. For this discussion point, we conducted a simplistic evaluation experiment that draws on the demonstrative examples presented in Sect. 5. It involves five artefacts using different versions of the HTML viewer presented in Sect. 5.1: the plain viewer, viewer that supports three types of visually impaired users (both versions presented in Sect. 5.1), viewer that dims the laptop’s screen when the battery’s charge goes below 50%<sup>9</sup>, and two additional equivalent implementations of the HTML viewer except with adaptations for the mentioned visually impaired support and dimming carried out by changing only the application level code (i.e., no change at the language level for the last two implementations) but still on the same variant of Neverlang.JS to keep the implementations comparable.

As a rough approximation for effort, we calculate the number of lines of application code as well as lines of language-level code needed for adaptations in each noted version. To minimise skill-dependency all examples were developed by an experienced Neverlang developer. To recap, we have a common basis (the plain HTML viewer) and two versions of comparable adaptations realised with and without micro-languages support. Table 3 shows the observed results in terms of lines of code (LOC) needed for these adaptations. The data for versions without/with micro-language support is shown on the left/right of the table respectively. The LOCs provided for the micro-languages-based adaptation include changes needed to modify the language and to adapt the application. The LOCs in the case without micro-languages implementation show changes needed to support the dynamic adaptation and those needed to implement the variants.

We observe that *there is a general reduction in the needed LOCs when micro-languages are used; the reduction is more evident at the application-level as expected.* The first

<sup>9</sup>This draws on the events used in the Mandelbrot example but acts on a different language facet to dim the screen when the laptop is unplugged.

row is for the accessibility adaptation; the second one for the dimming support. An average of 19 (26 and 12 respectively) LOCs was needed to provide a minimal structure for dynamic adaptivity without micro-languages. Moreover, about 215 LOCs were necessary to implement four different behaviours of the `print` language feature from Tab. 2. On the other hand, given an adaptivity framework as described in Fig. 3, adaptation through micro-languages would require no additional LOCs by the application programmer. Instead, the behavioural variants have to be implemented by the language programmer and we observe that about 57 LOCs were necessary to implement the statements from Tab. 2 and to program the adapter for micro-languages. The main difference is that with micro-languages the adaptation is confined to one point (the `print` statement implementation) and this is transparently adapted (we change the semantic of the statement not its syntax). Whereas without micro-languages, to smoothly support the change the `print` statement has been replaced by a call to the function implementing the behavior variability and to change each occurrence of the `print` statement represents an extra cost in terms of changed LOCs. The dimming adaptation is more contained in terms of modifications, we have only one variant instead of three as for long-sighted, colour-blind, and blind impediments. This explains the lower modified LOCs (64 for the application features and 48 for the language features). Please note that also using a pure Javascript implementation—i.e., without a `print` statement—for the adaptation at the application level the entity of the change would be of the same extent since we have still to adapt all the calls to the API used to `print`.

Thus, from our simplistic experiments, we conclude that our adaptation support has promise for reduction of code size and improved localisation of change, and modularity. It also proposes a new way of programming and composition, which cannot be accomplished without good tool support.

#### 6.4. Limitations and Open Challenges.

Several aspects of the proposed DSU approach are still challenging: i) the definition of the existing application features and their association to micro-languages, and ii) the use of program data—such as method and class names—to refine the area of influence of a language feature change.

Currently, which piece of code is part of an application feature is manually defined whereas the identification of the language features used to implement such a language feature—that is, the definition of the corresponding micro-language—is delegated to a variant of the language parser that recognized the used language constructs. A manual identification of the application features is a tedious and error-prone task and it is hardly maintainable; any change to the program will imply a reclassification of the code in application features that, if not done, would undermine the correctness of the  $\mu$ -DSU work. Automation of this task is, therefore, desirable and is part of on-going investigations. The easiest way should consist of asking the developer to decorate the code with meta-data about the piece of code that belongs to a particular application feature; similarly to what was done in [31, 32]. This approach would help with the maintenance issue but it still requires some manual identification with all the problems this entails. Some more automated approaches to concern mining are under consideration, such as associating an application feature with: each method as in [33], to all modules that calls the same method as in [34], code elements (methods, fields, ...) that present some naming similarities as in [35] or code fragments identical at the AST-level and widely reused as in [36]. All



these techniques are automated but focus on crosscutting application features whereas we need to identify any application feature.

Völter [37] defined the concepts of *linguistic* and *in-language abstractions* to denote the language elements—our language features—and program ones—class, field and method names and definitions—respectively. Both are abstractions that can represent update points.  $\mu$ -DSU focuses on linguistic abstractions but we are aware that some kind of adaptations—such as corrective adaptation—could be more easily achieved by addressing them at program level or at least having some in-language abstractions at disposal. As  $\mu$ -DSU is a dynamic software update framework, the program is available during its interpretation under the form of PT and the single program data (not the data used by it) as method names would be available through PT attributes. Therefore, in  $\mu$ -DSU, language adaptation is just the matter of either replacing a slice or injecting a semantic action at a specific node of the PT whereas a method change corresponds to a PT rewriting. Both approaches are feasible and supported by Neverlang as explained in [20] and [26] respectively. Nonetheless, we consider the support of changes at in-language abstractions too complex with respect to the support provided by other DSUs (that do not support changes at linguistic abstractions as  $\mu$ -DSU does) and we are still investigating how to limit this complexity before introducing in-language abstractions in  $\mu$ -DSU .

## 7. Related Work

Significant research has been done in the area of adaptive systems. The work here can be categorized as either *architectural* [38, 39] or *linguistic* [40, 41], or a combination of both [42]. Architectural approaches adapt an application by either adding, removing, or substituting one of its components. Traditional DSUs (such as JavAdaptor [1, 4], DUSC [5], Rubah [3] and JRebel [2]) are a variant of architectural approaches where evolution is not just a matter of system reconfiguration but also supports changes in the code. Such DSUs suffer from performance decay (due to indirections introduced by table forwarding and object proxies) [1, 5], limited program adaptation (no class re-positioning [2], either limited—UpgradeJ [43]—or no support for schema changes—HotSwap [44]), misalignment between design and executable code [31, 45] and a general difficulty in maintaining the evolved code [46, 47].

On the other hand, linguistic approaches support changes to both the application code and its behaviour. The idea to support evolution through ad hoc language constructs was first introduced in [48]. This idea is not orthogonal to the general DSUs approach, rather it allows to get the same evolution, with a better focus on the introduction of non-functional features. Today we have three types of mainstream linguistic approaches: *aspect-oriented programming* [49] (AOP), *reflection/meta-programming* [50] and *context-oriented programming* [51] (COP). Both AOP and reflection provide mechanisms to inject new code into an application yet keep the new code separate from the original one. AOP supports a complete decoupling of the new code (the crosscutting concerns) from the rest of the application code. Very few AO languages—such as CaesarJ [52] and AspectJ’s load-time weaving [27]—have any (limited) support for dynamic weaving at all, to update a running system. Despite AOP’s poor support for dynamic updating, a few DSUs were developed on top of it [53–55]. As an interesting by-product,  $\mu$ -DSU—and in particular the adaptation engine provided by Neverlang [20]—could be used as a

mechanism to implement dynamic aspect-oriented weaving at the language-level. Meta-programming approaches rely on the reflective features of the programming language and its runtime system. Reflection is used to observe and adapt the underlying program [56]. Chisel [57], PKUAS [58] and mChARM [59] are examples of frameworks that use Java’s reflection facilities to support software adaptation. COP introduces specific language-level abstractions for behavioural variations. Here the context becomes a first-class construct of a programming language [51, 60, 61]. The system then dynamically selects the appropriate behaviour or a combination of behaviours based on contextual information and selection conditions. The paradigm allows one to separate context-dependent crosscutting concerns. One drawback of this approach is that the computation and coordination aspects are often interleaved and behavioural variations must often be explicitly activated [62]. Moreover, the adaptation is implemented on a per-application basis. Mixins [63, 64] and traits [65] are other linguistic approaches to software adaptation and evolution; these approaches allow to extend a class with extra methods and override/enrich already present methods. Matrona [66] is a framework that uses mixins to support dynamic adaptation.

All these approaches enable solutions bound to the adaptation mechanisms provided by the language itself and they can hardly be ported to different programming languages and execution environments. They require adaptation to be explicitly implemented by the application developer in the application space itself. In support of this, Keeney *et al.* [57] argue that *“the software user, the developer, the designer and indeed the application logic itself all possess invaluable intelligence to gear how software should adapt itself to changing requirements and changing context”*. However, although highly valuable, we believe that this intelligence is rarely possessed by *all* the involved actors and thus it requires frequent collaboration between them, which goes against the principle of removing complexity from the lives of users and developers. On the other hand, as described in our approach, moving the adaptation to the level of micro-languages limits the complexity to a small group of developers which already possesses enough skills to handle the intricacy of complex systems. Furthermore, possible new behaviours affect a small subset of a language and the adaptation through micro-languages provides technical sustainability for free to all applications built upon them.

Kollár and Forgáč [67] investigated on the possibility of adapting a programming language interpreter during its execution. The proposed approach is based on AST rewriting—as implemented by Truffle [68]—and code injection at the PT nodes. Their proposal shares with  $\mu$ -DSU the idea of working on the PT but commonalities end there:  $\mu$ -DSU does not inject code but uses agents [20] that are dynamically hooked either before or after the PT node. Moreover,  $\mu$ -DSU confines the effect of the language change to a specific application feature thanks to the micro-languages and, thanks to  $\mu$ DA, it provides a easy way to specify the needed adaptation whereas the approach proposed by Kollár and Forgáč does not foresee any facility to deal with these aspects.

Research on modular language development and evolution is also related to our work. Languages evolve too [69, 70] and mechanisms to support their evolution are needed as well. Wassermann and Forgáč [71] survey all possible techniques to adapt domain specific languages implementation.  $\mu$ -DSU enables *grammar-based adaptation*—in Wassermann and Forgáč’s parlance—made dynamic thanks to the Neverlang’s parser, called DEXTER [25]. It permits to add and remove a grammar rule from an existing parser without regenerating it and therefore recompiling it. This is one of the key concepts to support language adaptation during the program execution. In the literature, it is

possible to find parsers with similar properties as Tatoo’s [72] and Silver’s [18] ones even if, in these cases, parser extensibility is limited to rules addition and not to their removal.

Modular language development frameworks although not directly related to software evolution, are enablers for micro-languages adaptation and as such they can represent a valid substitute to Neverlang in the application of the presented idea. Examples of modular language development frameworks are: Lisa [12], StrategoXT/Spoofax [11], Monticore [73], JastAdd [74], MPS [75] and Silver [18]. Some of these are comparatively analysed in [13, 76]. They all provide some mechanisms for modularising/composing a language, e.g., via inheritance in Lisa and Monticore, term rewriting in Spoofax and MPS and AST manipulation in JastAdd and MPS. Silver represents the closest approach to Neverlang, even if it is based on functional programming and pure attribute grammars. These characteristics enable all these frameworks to support software evolution via micro-languages. Yet, to our knowledge, none of them supports separate compilation and dynamic AST manipulation. Therefore any micro-language extension will imply recompiling the whole interpreter hindering the possibility of deploying adaptations at runtime. Tatoo [72, 77] focuses on grammars composition *de facto* enabling separate compilations but it does not support the modularization and composition of the language semantics: it could be integrated with any other modular language development framework to get the needed semantics support. More in general, any language development framework, such as Xtext [78], could provide micro-languages support to DSU, but modular development frameworks are preferable since they provide a finer change granularity and so are more flexible. Also approaches that enable language modularity through interpreters compositions—such as those approaches that use either *bridge interfaces* [79] or *language boxes* [80]—can be enablers for the micro-language support to DSU when the composed interpreters either have been developed by using modular language development framework or can be modularly modified.

## 8. Conclusions

This paper presents an approach to dynamic software updating through language semantics adaptation, using the novel concept of a micro-language. Essentially, micro-languages are a new, additional tool for modularising implementation of application features which: i) aligns the application features with the programming language implementing it; ii) separates the syntactic definition of control flow of the program from the semantics of the executed computation; iii) provides a new runtime composition mechanism for programs, thus expanding the toolset available for programming run-time adaptation. The approach has been implemented in the  $\mu$ -DSU framework on top of the Neverlang modular language development environment. Amongst the tasks that we will address in the immediate future are automation of application feature recovery along with identification of their associated micro-languages. This will extend the present work to deal with emergent and changing application features that would force the adaptation of the application. We will also develop a richer adaptation language, to substitute  $\mu$ DA, based on path temporal logic concepts. Further evaluation that also considers proficiency, security and efficiency aspects of this approach are currently underway.

## References

- [1] M. Pukall, C. Kästner, W. Cazzola, S. Götz, A. Grebhahn, R. Schöter, G. Saake, JavAdaptor — Flexible Runtime Updates of Java Applications, *Software—Practice and Experience* 43 (2) (2013) 153–185. doi:10.1002/spe.2107.
- [2] J. Kabanov, V. Vene, A Thousand Years of Productivity: The JRebel Story, *Software: Practice and Experience* 44 (1) (2014) 105–127.
- [3] L. Pina, L. Veiga, M. Hicks, Rubah: DSU for Java on a Stock JVM, in: T. Millstein (Ed.), *Proceedings of the ACM International Conference on Object Oriented Programming Systems Languages and Applications (OOPSLA'14)*, ACM, Portland, OR, USA, 2014, pp. 103–119.
- [4] M. Pukall, A. Grebhahn, R. Schröter, C. Kästner, W. Cazzola, S. Götz, JavAdaptor: Unrestricted Dynamic Software Updates for Java, in: *Proceedings of the 33<sup>rd</sup> International Conference on Software Engineering (ICSE'11)*, IEEE, Waikiki, Honolulu, Hawaii, 2011, pp. 989–991.
- [5] A. Orso, A. Rao, M. J. Harrold, A Technique for Dynamic Updating of Java Software, in: *Proceedings of the International Conference on Software Maintenance (ICSM'02)*, IEEE Press, Montréal, Canada, 2002, pp. 649–658.
- [6] S. Chiba, A Meta-Object Protocol for C++, in: *Proceedings of the 10th Annual Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'95)*, Vol. 30 of *Sigplan Notices*, ACM, Austin, Texas, USA, 1995, pp. 285–299.
- [7] A. Rashid, R. Chitchyan, Persistence As an Aspect, in: *Proceedings of the 2nd Int'l Conf. on Aspect-Oriented Software Development (AOSD'03)*, ACM Press, Boston, Massachusetts, 2003, pp. 120–129.
- [8] R. C. Gonçalves, J. L. Sobral, Pluggable Parallelisation, in: *Proceedings of the 18th ACM International Symposium on High Performance Distributed Computing (HPDC'09)*, ACM, Munich, Germany, 2009, pp. 11–20.
- [9] B. Harbulot, J. R. Gurd, Using AspectJ to Separate Concerns in Parallel Scientific Java Code, in: *Proceedings of the 3rd International Conference on Aspect-Oriented Software Development (AOSD'04)*, ACM Press, Lancaster, UK, 2004, pp. 122–131.
- [10] G. Bracha, Pluggable Type Systems, in: *Proceedings of ACM OOPSLA'04 Workshop on Revival of Dynamic Languages*, 2004.
- [11] L. C. L. Kats, E. Visser, The Spoofox Language Workbench: Rules for Declarative Specification of Languages and IDEs, in: M. Rinard, K. J. Sullivan, D. H. Steinberg (Eds.), *Proceedings of the ACM International Conference on Object-Oriented Programming Systems Languages and Applications (OOPSLA'10)*, ACM, Reno, Nevada, USA, 2010, pp. 444–463.
- [12] M. Mernik, An Object-Oriented Approach to Language Compositions for Software Language Engineering, *Journal of Systems and Software* 86 (9) (2013) 2451–2464.
- [13] E. Vacchi, W. Cazzola, Neverlang: A Framework for Feature-Oriented Language Development, *Computer Languages, Systems & Structures* 43 (3) (2015) 1–40. doi:10.1016/j.cl.2015.02.001.
- [14] R. Chitchyan, W. Cazzola, A. Rashid, Engineering Sustainability through Language, in: *Proceedings of the 37th International Conference on Software Engineering (ICSE'15)*, IEEE, Firenze, Italy, 2015, pp. 501–504, track on Software Engineering in Society.
- [15] W. Cazzola, D. M. Olivares, Gradually Learning Programming Supported by a Growable Programming Language, *IEEE Transactions on Emerging Topics in Computing* 4 (3) (2016) 404–415, special Issue on Emerging Trends in Education. doi:10.1109/TETC.2015.2446192.
- [16] K. Czarnecki, U. W. Eisenecker, *Generative Programming: Methods, Tools and Applications*, Addison-Wesley, 2000.
- [17] M. Bravenboer, K. T. Kalleberg, R. Vermaas, E. Visser, Stratego/XT 0.17. A Language and Toolset for Program Transformation, *Journal Science of Computer Programming* 72 (1-2) (2008) 52–70.
- [18] E. Van Wyk, D. Bodin, J. Gao, L. Krishnan, Silver: an Extensible Attribute Grammar System, *Electronic Notes on Theoretical Computer Science* 203 (2) (2008) 103–116.
- [19] W. Cazzola, A. Shaqiri, Dynamic Software Evolution through Interpreter Adaptation, in: *Proceedings of the 15th International Conference on Modularity (Modularity'16)*, ACM, Málaga, Spain, 2016, pp. 16–19.
- [20] W. Cazzola, A. Shaqiri, Open Programming Language Interpreters, *The Art, Science, and Engineering of Programming Journal* 1 (2) (2017) 5–1–5–34.
- [21] W. Cazzola, Domain-Specific Languages in Few Steps: The Neverlang Approach, in: T. Gschwind, F. De Paoli, V. Gruhn, M. Book (Eds.), *Proceedings of the 11<sup>th</sup> International Conference on Software Composition (SC'12)*, Lecture Notes in Computer Science 7306, Springer, Prague, Czech Republic, 2012, pp. 162–177.

- [22] W. Cazzola, E. Vacchi, Neverlang 2: Componentised Language Development for the JVM, in: W. Binder, E. Bodden, W. Löwe (Eds.), Proceedings of the 12<sup>th</sup> International Conference on Software Composition (SC'13), Lecture Notes in Computer Science 8088, Springer, Budapest, Hungary, 2013, pp. 17–32.
- [23] E. Vacchi, D. M. Olivares, A. Shaqiri, W. Cazzola, Neverlang 2: A Framework for Modular Language Implementation, in: Proceedings of the 13th International Conference on Modularity (Modularity'14), ACM, Lugano, Switzerland, 2014, pp. 23–26.
- [24] A. V. Aho, R. Sethi, J. D. Ullman, Compilers: Principles, Techniques, and Tools, Addison Wesley, Reading, Massachusetts, 1986.
- [25] W. Cazzola, E. Vacchi, On the Incremental Growth and Shrinkage of LR Goto-Graphs, *Acta Inf.* 51 (7) (2014) 419–447. doi:10.1007/s00236-014-0201-2.
- [26] W. Cazzola, A. Shaqiri, Modularity and Optimization in Synergy, in: D. Batory (Ed.), Proceedings of the 15th International Conference on Modularity (Modularity'16), ACM, Málaga, Spain, 2016, pp. 70–81.
- [27] G. Kiczales, E. Hilsdale, J. Hugunin, M. Kersten, J. Palm, B. Griswold, An Overview of AspectJ, in: J. L. Knudsen (Ed.), Proceedings of the 15th European Conference on Object-Oriented Programming (ECOOP'01), LNCS 2072, Springer-Verlag, Budapest, Hungary, 2001, pp. 327–353.
- [28] W. Cazzola, A. Cisternino, D. Colombo, Freely Annotating C#, *Journal of Object Technology* 4 (10) (2005) 31–48.
- [29] W. Cazzola, E. Vacchi, @Java: Bringing a Richer Annotation Model to Java, *Computer Languages, Systems & Structures* 40 (1) (2014) 2–18. doi:10.1016/j.cl.2014.02.002.
- [30] P. Greenwood, T. Bartolomei, E. Figueiredo, A. Garcia, N. Cacho, C. Sant'Anna, P. Borba, U. Kulesza, A. Rashid, On the Impact of Aspectual Decompositions on Design Stability: An Empirical Study, in: Proceedings of the 21<sup>st</sup> European Conference on Object-Oriented Programming (ECOOP'07), LNCS 4609, Springer-Verlag, Berlin, Germany, 2007, pp. 176–200.
- [31] W. Cazzola, S. Pini, A. Ghoneim, G. Saake, Co-Evolving Application Code and Design Models by Exploiting Meta-Data, in: Proceedings of the 22<sup>nd</sup> Annual ACM Symposium on Applied Computing (SAC'07), ACM Press, Seoul, South Korea, 2007, pp. 1275–1279.
- [32] M. Sulír, M. Nosál', J. Porubán, Recording Concerns in Source Code Using Annotations, *Computer Languages, Systems & Structures* 46 (2016) 44–65.
- [33] R. Serapilha Durelli, D. San Martín Santibáñez, N. Anquetil, M. E. Delamaro, V. Vieira de Camargo, A Systematic Review on Mining Techniques for Crosscutting Concerns, in: Proceedings of the 28th Annual ACM Symposium on Applied Computing (SAC'13), ACM Press, Coimbra, Portugal, 2013, pp. 1080–1087.
- [34] D. Zhang, Y. Guo, X. Chen, Automated Aspect Recommendation through Clustering-Based Fan-In Analysis, in: A. Ireland, W. Visser (Eds.), Proceedings of the 23rd IEEE/ACM International Conference on Automated Software Engineering (ASE'08), IEEE, L'Aquila, Italy, 2008, pp. 278–287.
- [35] D. Shepherd, L. Pollock, T. Tourwé, Using Language Clues to Discover Crosscutting Concerns, in: Proceedings of the Workshop on Modeling and Analysis of Concerns in Software (MACS'05), IEEE, St. Louis, Missouri, USA, 2005, pp. 1–6.
- [36] M. Bruntink, A. van Deursen, R. van Engelen, T. Tourwé, On the Use of Clone Detection for Identifying Crosscutting Concern Code, *IEEE Transactions on Software Engineering* 31 (10) (2005) 804–818.
- [37] M. Voelker, *DSL Engineering*, 2013.
- [38] P. Oreizy, N. Medvidovic, R. N. Taylor, Architecture-Based Runtime Software Evolution, in: Proceedings of the 20<sup>th</sup> International Conference on Software Engineering (ICSE'98), IEEE Computer Society, Kyoto, Japan, 1998, pp. 177–186.
- [39] J. Kramer, J. Magee, Self-Managed Systems: an Architectural Challenge, in: L. C. Briand, A. L. Wolf (Eds.), Proceedings of 29<sup>th</sup> International Conference on Software Engineering (ICSE'07): Future of Software Engineering (FoSE'07), IEEE Computer Society, Minneapolis, MN, USA, 2007, pp. 259–268.
- [40] G. Salvaneschi, C. Ghezzi, M. Pradella, An Analysis of Language-Level Support for Self-Adaptive Software, *ACM Transactions on Autonomous and Adaptive Systems* 8 (2) (2013) 7:1–7:29.
- [41] L. Galletta, *Adaptivity: Linguistic Mechanisms and Static Analysis Techniques*, PhD Thesis, Università degli Studi di Pisa, Pisa, Italy (May 2014).
- [42] M. Salehie, L. Tahvildari, Self-Adaptive Software: Landscape and Research Challenges, *ACM Transactions on Autonomous and Adaptive Systems* 4 (2) (2009) 14:1–14:42.
- [43] G. Bierman, M. Parkinson, J. Noble, UpgradeJ: Incremental Typechecking for Class Upgrades, in: J. Vitek (Ed.), Proceedings of the 22nd European Conference on Object-Oriented Programming

- (ECOOP'08), Lecture Notes in Computer Science 5142, Springer, Paphos, Cyprus, 2008, pp. 235–259.
- [44] M. Dmitriev, Towards Flexible and Safe Technology for Runtime Evolution of Java Language Applications, in: V. Cahill, S. Clarke, S. Dobson, R. Filman (Eds.), Proceedings of the 1st Workshop on Engineering Complex Object-Oriented Systems for Evolution (ECOOSE'01), Tampa Bay, FL, USA, 2001, pp. 14–18.
- [45] G. C. Murphy, D. Notkin, K. J. Sullivan, Software Reflexion Models: Bridging the Gap between Design and Implementation, *IEEE Transactions on Software Engineering* 27 (4) (2001) 364–380.
- [46] T. D'Hondt, K. De Volder, K. Mens, R. Wuyts, Co-Evolution of Object-Oriented Software Design and Implementation, in: M. Akşit (Ed.), Proceedings of the International Symposium on Software Architectures and Component Technology, Kluwer, Twente, The Netherlands, 2000, pp. 207–224.
- [47] P. Ebraert, Y. Vandewoude, T. D'Hont, Y. Berbers, Pitfalls in Unanticipated Dynamic Software Evolution, in: W. Cazzola, S. Chiba, G. Saake, T. Tourwé (Eds.), Proceedings of ECOOP'2005 Workshop on Reflection, AOP and Meta-Data for Software Evolution (RAM-SE'05), Glasgow, Scotland, 2005, pp. 3–8.
- [48] T. Mens, M. Wermelinger, S. Ducasse, S. Demeyer, R. Hirschfeld, M. Jazayeri, Challenges in Software Evolution, in: Proceedings of the Eighth International Workshop on Principles of Software Evolution (IWPSE'05), IEEE Press, Lisbon, Portugal, 2005, pp. 13–22.
- [49] G. Kiczales, J. Lamping, A. Mendhekar, C. Maeda, C. Videira Lopes, J.-M. Loingtier, J. Irwin, Aspect-Oriented Programming, in: 11th European Conference on Object Oriented Programming (ECOOP'97), Lecture Notes in Computer Science 1241, Springer-Verlag, Helsinki, Finland, 1997, pp. 220–242.
- [50] P. Maes, Concepts and Experiments in Computational Reflection, in: N. K. Meyrowitz (Ed.), Proceedings of the 2nd Conference on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA'87), Vol. 22 of Sigplan Notices, ACM, Orlando, Florida, USA, 1987, pp. 147–156.
- [51] R. Hirschfeld, P. Costanza, O. Nierstrasz, Context-Oriented Programming, *Journal of Object Technology* 7 (3) (2008) 125–151.
- [52] I. Aracic, V. Gasiunas, M. Mezini, K. Ostermann, An Overview of CaesarJ, *Transaction on Aspect-Oriented Software Development* 1 (1) (2006) 135–173.
- [53] Z. Yang, B. H. C. Cheng, R. E. K. Stirewalt, J. Sowell, S. M. Sadjadi, P. K. McKinley, An Aspect-Oriented Approach to Dynamic Adaptation, in: D. Garlan, J. Kramer, A. L. Wolf (Eds.), Proceedings of the 1st Workshop on Self-Healing Systems (WOSS'02), ACM, Charleston, SC, USA, 2002, pp. 85–92.
- [54] P. Greenwood, L. Blair, A Framework for Policy Driven Auto-adaptive Systems Using Dynamic Framed Aspects, *Transactions on Aspect-Oriented Software Development* 2 (2006) 30–65.
- [55] G. Zhang, M. Rong, A Framework for Dynamic Evolution Based on Reflective Aspect-Oriented Software Architecture, in: Proceedings of the 4<sup>th</sup> International Conference on Computer Sciences and Convergence Information Technology (ICCIT'09), Seoul, South Korea, 2009, pp. 7–10.
- [56] O. Nierstrasz, T. Girba, Lessons in Software Evolution Learned by Listening to Smalltalk, in: J. Leeuwen, A. Muscholl, D. Peleg, J. Pokorný, B. Rumpe (Eds.), Proceedings of the 36th Conference on Current Trends in Theory and Practice of Computer Science (SOFSEM'10), LNCS 5901, Springer, Špindlerův Mlýn, Czech Republic, 2010, pp. 77–95.
- [57] J. Keeney, V. Cahill, Chisel: A Policy-Driven, Context-Aware, Dynamic Adaptation Framework, in: Proceedings of the 4th International Workshop on Policies for Distributed Systems and Networks (POLICY'03), IEEE, Como, Italy, 2003, pp. 3–14.
- [58] G. Huang, H. Mei, F.-Q. Yang, Runtime Software Architecture Based on Reflective Middleware, *Journal of Information Science* 47 (5) (2004) 555–576.
- [59] W. Cazzola, Remote Method Invocation as a First-Class Citizen, *Distributed Computing* 16 (4) (2003) 287–306. doi:10.1007/s00446-003-0094-8.
- [60] R. Keays, A. Rakotonirainy, Context-Oriented Programming, in: S. Banerjee, M. Cherniack (Eds.), Proceedings of the 3rd ACM International Workshop on Data Engineering for Wireless and Mobile Access (MobiDe'03), ACM, San Diego, CA, USA, 2003, pp. 9–16.
- [61] S. González, K. Mens, M. Colăcioiu, W. Cazzola, Context Traits: Dynamic Behaviour Adaptation through Run-Time Trait Recomposition, in: J. Kienzle (Ed.), Proceedings of the 12th International Conference on Aspect-Oriented Software Development (AOSD'13), ACM, Fukuoka, Japan, 2013, pp. 209–220.
- [62] G. Salvaneschi, C. Ghezzi, M. Pradella, Context-Oriented Programming: A Software Engineering Perspective, *Journal of Systems and Software* 85 (8) (2012) 1801–1817.
- [63] G. Bracha, W. Cook, Mixin-Based Inheritance, in: A. Yonezawa (Ed.), Proceedings of the European

- Conference on Object-Oriented Programming on Object-Oriented Programming Systems, Languages, and Applications (OOPSLA/ECOOP'90), ACM, Ottawa, Canada, 1990, pp. 303–311.
- [64] M. Flatt, S. Krishnamurthi, M. Felleisen, Classes and Mixins, in: D. B. MacQueen, L. Cardelli (Eds.), Proceedings of the 25th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (PoPL'98), ACM, 1998, pp. 171–183.
- [65] N. Schärli, S. Ducasse, O. Nierstrasz, A. P. Black, Traits: Composable Units of Behaviour, in: L. Cardelli (Ed.), Proceedings of the 17<sup>th</sup> European Conference on Object-Oriented Programming (ECOOP'03), Lecture Notes in Computer Science 2743, Springer, Darmstadt, Germany, 2003, pp. 248–274.
- [66] M. Springer, F. Niephaus, R. Hirschfeld, H. Masuhara, Matriona: Class Nesting with Parametrization in Squeak/Smalltalk, in: D. Batory (Ed.), Proceedings of the 15th International Conference on Modularity (Modularity'16), ACM, Málaga, Spain, 2016, pp. 118–129.
- [67] J. Kollár, M. Forgáč, Combined Approach to Program and Language Evolution, *Computing and Informatics* 29 (6) (2010) 1103–1116.
- [68] T. Würthinger, A. Wöß, L. Stadler, G. Duboscq, D. Simon, C. Wimmer, Self-Optimizing AST Interpreters, in: A. Warth (Ed.), Proceedings of the 8<sup>th</sup> Symposium on Dynamic Languages (DSL'12), ACM, Tucson, AZ, USA, 2012, pp. 73–82.
- [69] J.-M. Favre, Languages Evolve Too! Changing the Software Time Scale, in: M. Saeki, G. Canfora, S.-i. Yamamoto (Eds.), Proceedings of the 8th International Workshop on Principles of Software Evolution (IWPE'05), Lisbon, Portugal, 2005, pp. 33–42.
- [70] W. Cazzola, D. Poletti, DSL Evolution through Composition, in: Proceedings of the 7<sup>th</sup> ECOOP Workshop on Reflection, AOP and Meta-Data for Software Evolution (RAM-SE'10), ACM, Maribor, Slovenia, 2010.
- [71] Ľ. Wassermann, M. Forgáč, Adaptation Techniques of Domain-Specific Languages, in: Proceedings of the 10th Scientific Conference of Young Researchers (SCYR'10), FEI TU of Košice, Košice, Slovakia, 2010, pp. 370–373.
- [72] J. Cervelle, R. Forax, G. Roussel, A Simple Implementation of Grammar Libraries, *Computer Science and Information Systems* 4 (2) (2007) 65–77.
- [73] H. Krahn, B. Rumpe, S. Völkel, MontiCore: A Framework for Compositional Development of Domain Specific Languages, *International Journal on Software Tools for Technology Transfer* 12 (5) (2010) 353–372.
- [74] G. Hedin, E. Magnusson, JastAdd — An Aspect-Oriented Compiler Construction System, *Science of Computer Programming* 47 (1) (2003) 37–58.
- [75] M. Völter, V. Pech, Language Modularity with the MPS Language Workbench, in: Proceedings of the 34th International Conference on Software Engineering (ICSE'12), IEEE, Zürich, Switzerland, 2012, pp. 1449–1450.
- [76] S. Erdweg, T. van der Storm, M. Völter, L. Tratt, R. Bosman, W. R. Cook, A. Gerritsen, A. Hulshout, A. Kelly, G. Konat, P. J. Molina, M. Palatnik, R. Pohjonen, E. Schindler, K. Schindler, R. Solmi, V. Vergu, E. Visser, K. van der Vlist, G. Wachsmuth, J. van der Woning, Evaluating and Comparing Language Workbenches: Existing Results and Benchmarks for the Future, *Computer Languages, Systems and Structures* 44 (2015) 24–47.
- [77] J. Cervelle, R. Forax, G. Roussel, Tatoo: An Innovative Parser Generator, in: R. Gitzel, M. Aleksey, M. Shader (Eds.), Proceedings of the 4th International Symposium on Principles and Practice of Programming in Java (PPPJ'06), Mannheim, Germany, 2006, pp. 13–20.
- [78] S. Efftinge, M. Völter, oAW xText: A Framework for Textual DSLs, in: Proceedings of the EclipseCon Summit Europe 2006 (ESE'06), Vol. 32, Esslingen, Germany, 2006.
- [79] M. Grimmer, C. Seaton, T. Würthinger, H. Mössenböck, Dynamically Composing Languages in a Modular Way: Supporting C Extensions for Dynamic Languages, in: G. Leavens (Ed.), Proceedings of the 14th International Conference on Modularity (Modularity'15), ACM, Fort Collins, CO, USA, 2015, pp. 1–13.
- [80] E. Barrett, C. F. Bolz, L. Tratt, Approaches to Interpreter Composition, *Computer Languages, Systems & Structures* 44 (Part C) (2015) 199–217.