



From Specification to Testing: Semantics Engineering for Lua

5.2

Mallku Soldevila¹ · Beta Ziliani² · Bruno Silvestre³

Received: 6 July 2021 / Accepted: 22 June 2022 / Published online: 11 August 2022
© The Author(s), under exclusive licence to Springer Nature B.V. 2022

Abstract

We provide a formal semantics for a large subset of the Lua programming language, in its version 5.2. The semantics is a major part of an ongoing effort to construct reliable tools to analyze Lua code. In this work, we present the details of several key aspects of the language, like the semantics of its only structured data-type (*tables*), its meta-programming mechanism (*metatables*), error handling, and how these mechanisms are used to define a complex dynamic semantics that must deal with several possible erroneous situations during run time, given the nature of the language. The semantics is mechanized in Redex, a DSL specially designed to specify and debug operational semantics. We validated the mechanization in two ways: first, by executing within Redex the test suite of the reference interpreter of Lua, and second, by specifying and performing random testing of its fundamental properties using the redex-check tool. Together, they evidence that our model soundly captures the semantics of the selected fragment of the language. Additionally, we address some of the performance problems that typically arise when testing a mechanization in Redex, by using a simple implementation of a reachability-based garbage collector that captures key aspects of Lua's. By collecting syntactic garbage, we reduce the size of configurations during run time. Finally, we briefly discuss this avenue of development of our semantics, together with the implementation of a prototype tool to perform static analysis of Lua programs.

Keywords Semantics · Operational semantics · Imperative languages · Domain specific languages · Lua · Reduction semantics · Randomized testing

Mathematics Subject Classification 10011311 · 10010134 · 10011010 · 10011017

✉ Mallku Soldevila
mes0107@famaf.unc.edu.ar

Beta Ziliani
beta@mpi-sws.org

Bruno Silvestre
brunoos@inf.ufg.br

¹ FAMAF, UNC and CONICET, Córdoba, Argentina

² FAMAF, UNC and Manas.Tech, Córdoba, Argentina

³ INF, UFG, Goiânia, Brazil

1 Introduction

Lua is a lightweight imperative scripting language, featuring dynamic typing, automatic memory management, data description facilities, and meta-programming mechanisms to adapt the language to specific domains [11]. The typical use case of a Lua application is as an extension library embedded in a *host* application, commonly written in C or C++. In that setting, Lua offers the possibility to add scripting facilities to the host application, combining the flexibility and rapid prototyping of a dynamic language within the static guarantees and optimizations of stricter programming languages.

Lua is extensively used in many diverse applications, ranging from game development, most notably by “AAA” games [12] but also in mobile games and game frameworks, plug-in development (for example, in the photo-editing software Adobe Photoshop Lightroom [1] and the type-setting system LuaTeX [28]), web application firewalls [7], and embedded systems [27].

Lua is informally specified by both its reference manual and its reference interpreter, developed and maintained by the core Lua authors. Thanks to Lua’s success, several alternative implementations as well as code linters and static analyzers can be found in the wild [24, 25]. However, the informal nature of the specification means that developers of these tools must resort to their intuition, formed by study of the reference manual, inspection of the source code of the interpreter, and experimentation.

Even if Lua is considered a small language, it is not unusual to stumble across some puzzling behaviors that may not be easily inferred from the official documentation. As a simple example consider the following program:

Example 1 (To return or not to return.)

```

1 function f() return nil end
2
3 function g() end
4
5 print(f(), f(), f()) -->> nil nil nil
6 print(g(), g(), g()) -->> nil nil

```

In this example, we have *f* returning Lua’s null value (**nil**), and *g* not returning anything. Then, when printing three identical calls to *f* and *g*, respectively, we see a difference: in the former case three **nil**s are printed, while in the latter only two are (we write `-->>` to indicate the output of the interpreter; `--` is Lua’s token for comments). The different outputs can be explained by two related aspects of Lua: *tuples* and *vararg* arguments. Briefly, some Lua functions can receive a list of arbitrary length of actual parameters. Such functions are known as *vararg functions*. The **print** service is an example of a *vararg* function. Internally, the list of zero or more actual parameters are manipulated through what we will call a *tuple*. Also, a function can return a list of zero or more values, which will be also modeled as tuples. Over a tuple there are a set of rules that apply, in order to extract their values and use them in the context where the tuple appears: if it appears into a list of values (for example, the list of actual parameters in a function call, as in lines 5–6 in the code shown), we need to append the tuple’s values to the list; if it appears where a single value is required, we need to *truncate* the tuple and extract just its first value. As in the example shown, said rules may shield unexpected results, mostly when we deal with *empty* tuples.

In this work, we present an extensive formalization—mechanized in Redex [5]—of a large subset of Lua 5.2, explaining in detail its different mechanisms, like its particular take on

tuples, and providing sound grounds for developing alternative implementations, extensions, and analysis tools for the language. In fact, the present work refines, extends, and greatly improves [37], which was the basis to formalize Lua's garbage collector [38], including the development of a small tool to verify simple programs making use of Lua's *weak tables* (akin to weak references). As a side note, Lua 5.2 includes interfaces with the garbage collector, namely *weak tables* and *finalizers*, which have an observable impact on the behavior of user programs. In order to account for this, we include on top of our semantics a simple syntactic, non-deterministic, garbage collector to which it is possible to communicate through said interfaces, and we show how the impact on user programs' semantics manifests in that setting.

In order to gain confidence in our model, we took three measures: first, we developed a semantics that (mostly) contains each programming concept from the reference manual, in a one-to-one correspondence, instead of focusing on an idealized core of the language. Second, we successfully run the test suite of the reference interpreter of the language [21] *directly within Redex*, checking that our semantics is correct to the extent of what is tested within the suite. That is, unlike many other works in the literature (e.g., [8, 35]), we did not pay the cost—confidence-wise—of building a new interpreter. Last, we used Redex's support for formal systems and random testing [15] to random test soundness properties of our semantics. This allowed us to detect a plethora of errors and omissions that were not caught during the development of the semantics, even after successfully passing the tests of the test suite.

In overall, we obtain evidence that our semantics is sound and corresponds to that of the selected subset of the language's features, including

- Every type of Lua value, except *coroutines* and *userdata* (see below);
- Tuples;
- Metatables (Lua's configurable dynamic dispatch);
- Identity of closures;
- Dynamic execution of source code;
- Error handling;
- A large collection of services from the standard library;
- Garbage collection (presented in [38]).

For reasons of simplicity, we purposely excluded the following features for future work:

- Coroutines, in essence single-shot delimited continuations;
- Userdata, opaque handles to data from the host application and native libraries;
- The **goto** statement;
- Services from the standard library that interface with the operating system, such as file manipulation, or have large complex C implementations, such as string pattern matching.

We develop further the main reasons behind the exclusion of some of these features in Sect. 9.

1.1 Contributions

As a summary, we present

1. A formalization of a large portion of Lua, in its version 5.2, including several semantics relevant details of the official interpreter not covered in the reference manual.
2. A mechanization of said formalization in Redex.

3. An interpreter based on the mechanization that successfully passes all relevant tests cases from Lua's interpreter test suite.
4. A formalization and a randomly tested mechanization of the soundness property of the semantics.
5. A brief discussion of applications of the semantics model presented.

The mechanization can be downloaded from

<https://github.com/Mallku2/lua-gc-redex-model>

1.2 Quick Tour

Section 2 presents a brief description of Lua, with emphasis on particular features that we formalize in later sections. Its reading is non-essential for the Lua *connoisseur*.

Section 3 presents the basic concepts that our model uses, via a formalization of a very small subset of Lua. Its reading is important to understand the following sections.

Example 2 (Memoization in Lua.)

```

1 local function memoize(fn)
2   local t = {}
3   return function(x)
4     local y = t[x]
5     if y == nil then y = fn(x); t[x] = y end
6     return y
7   end
8 end
9 local memsum = memoize(function(x)
10  local a = 1
11  for i = 1,x do a = a + i end
12  return a
13 end)

```

Section 4 expands Sect. 3 with the formalization of the most interesting parts of the semantics. Of special interests are the formalization of metatables (Sect. 4.4) and tuples (Sect. 4.2.3).

Section 5 presents some relevant properties of the language.

Section 6 presents the software artifact (the mechanization in Redex of the language), and the results of the two types of testing performed. It is an essential part of the work; its reading being most recommended.

Section 7 introduces two related applications of this work, modeling garbage collection, and developing static analysis tools.

Section 8 discusses related work.

Section 9 reflects on the experience of formalizing and mechanizing Lua, and discusses future avenues of research.

2 Lua: An Extensible Scripting Language

In this section, we present a quick tour to some of Lua's most salient characteristics and features captured in our model. The Lua expert can safely skip this section and come back to it only if necessary.

2.1 First-Class Closures and Tables

We start with a simple memoization function,¹ listed in Example 2, which takes a function `fn` as argument and returns its *memoized* version. The values of `fn` already computed will be stored in a *table* (`t` in line 2). Tables are, in essence, associative arrays indexed by any Lua value except the null value (**nil**). They also come with syntax sugar and metaprogramming facilities that can greatly extend their functionality beyond simple associative arrays, as we will see in Sect. 2.3 below.

Example 3 (The environment `_ENV`)

```

1 x = 10
2 print(_ENV['x'])  -->> 10
3 _ENV['x'] = 0
4 print(x)  -->> 0
5 _ENV = {}
6 print()  -->> error, print is not anymore in _ENV

```

The memoized version of `fn` is returned through an anonymous function in line 3. This function takes `x` as argument and, before computing `fn(x)`, performs a look-up in the table for value `x` (line 4). If the result of the look-up is **nil**, it means no result was found, so it proceeds to compute `fn(x)` and store it in the table (line 5). The resulting value, either computed or retrieved from the table, is returned in line 6. Note that the memoization only works for values `fn(x)` different than **nil**. The function `memoize` is used in lines 10–14 to improve the performance of a function that performs a sum from 1 to `x`.

It is important to note that all procedures in Lua, anonymous or named, are first-class values, and form lexically scoped closures. The anonymous function that `memoize` returns will effectively capture into its scope the table `t`, as expected.

2.2 Local Definitions and the Environment `_ENV`

Note that the definitions of `memoize`, `t`, and `memsum` are prefixed by the keyword **local**. As its name suggests, **local** creates a local variable, initialized with the *rvalue* (what comes after the '=' operator). If we omit this keyword, a declaration is considered an *assignment*, and if the variable was declared local before, then it performs an update of its value, as usual. More interestingly, if there is no such variable in scope, then Lua looks up for the variable in a table called the *environment*, which is bound to the identifier `_ENV`, using the variable's identifier as the key. Similarly, reading the value of a variable that is not within scope triggers the search for the value associated with that name in the environment. Abusing the notation, this means that any occurrence of a variable `x` that is *not* in scope, is just syntax sugar for `_ENV["x"]`. Variables in the environment are called *global* variables. Services from the standard library (like `print`) are accessible through global variables, and are, therefore, bound by this environment.

Since `_ENV` is in itself a variable, the programmer can, at any time, change the environment in which a program is executed by simply assigning another table to `_ENV`. Example 3 shows some simple interactions with the environment to illustrate this concept.

¹ Taken from <http://lua-users.org/wiki/FuncTables>.

2.3 Metatables

The most notable feature of Lua is its meta-programming mechanism, *metatables*, that lets the programmer adapt the language to specific domains. Thanks to metatables, Lua can maintain its original design decision to *keep the language simple and small* [12], while still being able to cope with a variety of programming concepts (Section “Code Structure/Programming Paradigms” of [26]).

Example 4 (OOP based on Lua’s metatable mechanism.)

```

1 local MyClass = {}
2 MyClass.__index = MyClass
3
4 function MyClass.new(init)
5   local self = setmetatable({}, MyClass)
6   self.value = init
7   return self
8 end
9
10 function MyClass:set_value(newval)
11   self.value = newval
12 end
13
14 function MyClass:get_value()
15   return self.value
16 end
17
18 local mc = MyClass.new(5)
19 print(mc:get_value()) -->> 5
20 mc:set_value(6)
21 print(mc:get_value()) -->> 6

```

As an example, we show the implementation of some basic concepts of object-oriented programming. The example² is listed in Example 4, and it models classes and objects by combining tables, first-class functions—which we already saw in Sect. 2.1—and the metatable mechanism. It also introduces some syntax sugar provided by Lua to better support OOP.

In Lua, a class is essentially implemented as a dictionary (e.g., table), in which the method names form the keys of the dictionary, and the method implementations are the associated values. Objects are also modeled with tables, containing in their fields their attributes.

In the example, we have a class `MyClass` with its corresponding constructor (line 4) and only one field `value` with its setter (line 10) and getter (line 14). The function declarations in these lines are actually syntax sugar for assignments, where the left-hand sides are, respectively, `MyClass["new"]`, `MyClass["set_value"]`, and `MyClass["get_value"]`. For the two methods on line 10 and line 14, the use of `:` instead of `.` also includes an implicit first parameter for these functions, named `self`.

In the last lines of Fig. 4, we show how to create an instance of `MyClass` (line 18), and how to invoke the methods. In line 20, we can observe the invocation of `set_value` with yet another syntax sugar: `mc:set_value(6)` is equivalent to `mc["set_value"](mc, 6)` (remember that the first, implicit, parameter of `set_value` is `self`).

² Taken from <http://lua-users.org/wiki/ObjectOrientationTutorial>.

Example 5 (Taking the maximum element of an array, together with its index.)

```

1 function maximum (a)
2   local mi = 1           -- maximum index
3   local m = a[mi]       -- maximum value
4   for i, val in ipairs(a) do
5     if val > m then
6       mi = i
7       m = val
8     end
9   end
10  return m, mi
11 end
12
13 print(maximum({8,10,23,12,5})) -->> 23 3

```

If classes contain methods, and objects contain fields, how is `mc["set_value"]` looking up the `set_value` method? The answer is the *metatable* mechanism, used in lines 2 and 5. In line 5, the call to `setmetatable` assigns `MyClass` as the metatable of the table provided as argument (the empty table `{}`). The (empty) table returned is our self object.

A metatable can modify the behavior of a table with regard to most of Lua's operations. For this example, the behavior we are modifying is look-up of non-existing keys. Each behavior can be modified through a *metamethod* that, despite its name, is a *value* associated with a specific key in the metatable. In this case, the key is `__index` (line 2), and the value is the table `MyClass`. The overall effect will be that a miss when looking up for a key in the table (`mc`) will trigger a look-up for the same key in the metatable (`MyClass`). This is how `mc["set_value"]` results in the method `set_value` from `MyClass`.

Lua also allows metamethods for an indexing with a non-existing key, for a function call over a value that is not a closure, for unexpected circumstances involving binary and unary operators, for setting finalizers, and even for some functions in the standard library. Lua programmers typically use metatables for object-oriented programming (including more elaborated object models than class-based single inheritance), for operator overloading, and for proxies.

2.4 Tuples

A Lua function may return several values in what we call a *tuple*. Tuples are particular in the sense that several of Lua constructs treat them specially. Take for instance Example 5 ([10], Chp. 5.1). The function `maximum` takes an array `a` and returns its maximum element, together with its index. It uses the *iterator*, obtained from the function `ipairs` of the standard library (line 4), which, in a first call, returns a tuple with three values: an iterator function, the table to be iterated, and 0. By calling repeatedly this iterator, supplying the table to be iterated and a position, we can traverse the table following its numeric integer indexes, beginning at 1. Accordingly, the **for-in** construct knows how to use the iterator to obtain, at each iteration, an index to the table being iterated and the value at that position. When the for-loop ends, we have in `m` the maximum value, and in `mi`, its index in the array. These are returned by `maximum` in line 10.

Example 6 (Converting an array of elements into a tuple.)

```

1 function unpack (t, i)

```

```

2  i = i or 1
3  if t[i] ~= nil then
4    return t[i], unpack(t, i + 1)
5  end
6 end
7
8 a1, a2 = unpack ({1,2,3})  -- a1 = 1 and a2 = 2
9
10 b1, b2, b3, b4 = unpack ({1,2,3})  -- b1 = 1, b2 = 2, b3 = 3, b4 = nil

```

The size of a tuple does not need to be statically determined. An example of dynamically constructed tuples is presented in Example 6, where function `unpack` takes an array of elements and produces a tuple containing such elements. There are a couple of niceties in the code that needs some explanation. To begin with, in Lua, a function can be called with fewer or extra arguments. Every argument that is not passed by the caller of the function has value `nil`, and every extra argument is simply discarded. This feature allows, for example, to avoid having an auxiliary function performing recursion, and then the main function calling it with the base case. Instead, we can combine the two. In `unpack`, the index `i` used to index the array can be omitted, in which case we ensure it starts at 1 (arrays are 1-based indexed) thanks to the `or` operator (line 2). If `i` is `nil`, then `or` will return its second argument, 1.

Another aspect of tuples worth discussing is that we can choose to ignore elements from a tuple, or even *take* more elements than actually present in the tuple. The first case is illustrated in line 8, where we assign only the first two elements from the tuple returned by `unpack`, discarding the third. The second is illustrated in line 10, where the first three variables get the three elements from the tuple, while the fourth gets value `nil`.

2.5 Vararg Arguments

It is possible to specify a *variadic* function, *i.e.*, a function that can receive an arbitrary quantity of parameters, by appending the *vararg* parameter to the list of formal parameters of a function. Noted with `...`, this argument collects all the extra arguments that are provided to a function. In Fig. 7, we see a function dual to `unpack`: it receives a list of arguments of arbitrary length and returns a *vector* containing those arguments: a table whose elements are indexed by natural numbers.

Example 7 (Converting a list of arguments into a table.)

```

1 function pack (...)
2   return {...}
3 end
4
5 t = pack (1,2,3)
6 print(t [1], t [2], t [3], t [4])  -->> 1  2  3  nil

```

3 Basics of the Formalization

In this section, we gently introduce the semantic framework used throughout the paper by providing semantics to Lua_0 , a toy subset of Lua. We write the semantics in a small-step operational semantics with evaluation contexts. Evaluation contexts, taken from Felleisen–Hieb’s Reduction Semantics (RS) [5], are used for the specific purposes of modularization;

for providing a concise description of the context-sensitive semantics; and to define the execution order. In this, we follow the path taken by [8, 34, 35], where RS is similarly applied for successfully formalizing real programming languages. However, we choose a different path with regard to the presentation of our formal semantics and the relation with its mechanization:

1. We emphasize the distinction between what constitutes the language we want to model, and what are the *run-time constructs*: concepts proper of the dynamic semantics required to provide an operational semantics.
2. We reduce the complexity of the *desugaring* process, by staying as close as possible to the source language.
3. We maintain a short-yet-noticeable distance between what is the formal model presented in the paper and its corresponding Redex mechanization.

The first two respond to the need of building the model as comprehensible as possible, mainly for the developers of the language and tools: Point 1 helps distinguishing the syntax developers care about from the constructs that are not visible to them when executing programs, and Point 2 reduces the gap between the code shown in the rules and actual Lua code, which helps building *trust* in the model (also noted in [2, 29]).

The motivation behind Point 3—discussed in depth in Sect. 6—is to present the rules of the language in a more natural form (as seen in many books on semantics, e.g., [33]), instead of presenting Redex code directly.

For the small subset of the Lua language presented in this section, we have two fragments: *pure* statements and expressions (following Lua’s distinction of statements and expressions), and *stateful* (i.e., memory changing) statements. We describe each of them in isolation, each with their own *relation*. Then, we compose the two using a third relation which will also deal with execution order, providing semantics to entire programs.

Some comments on notation: in the following, we will use *nt* to refer to, either, the non-terminal *nt* or the set of terms generated from the production of said non-terminal (resorting to context for disambiguation); an *nt*, possibly with a numeric sub-index, to refer to a term that belongs to the set of terms generated by *nt*.

3.1 The Pure Fragment of Lua₀

We show the grammar and semantics for *stateless* programs in Fig. 1. As statements (*s*), we only include conditional branching and skip (*;*). The condition is an expression *e*, which can be a *value* (*v*), a fully applied and infix binary operator (*binop*), or a fully applied and prefix unary operator (*unop*). Values are **nil** or boolean literals (**true** or **false**). Operators are the logical connectives **and**, **or**, and **not**. Of course, with this language, we are not able to write any useful program, but in the coming sections, we will grow the language until we reach Lua.

The operational semantics for the pure fragment are modeled with the $\rightarrow^{s/e}$ relation between statements and expressions. Rule IF- T states that if the conditional of the **if** is any value different from **nil** and **false**, then it is considered **true**, and therefore, the **then** branch is taken. Note that we write above the line the conditions in which the rule applies. We omit the line in the case that no condition is required. Rule IF- F states that, for **false** or **nil**, the **else** branch is taken.

The reader may wonder why we have not put **true** instead of *v* in the first rule, since there are no other values. The reason is that we will grow the language to have more values, like

Grammar

$s ::= \text{if } e \text{ then } s \text{ else } s \mid ;$ $\text{binop} ::= \text{and} \mid \text{or}$
 $e ::= v \mid e \text{ binop } e \mid \text{unop } e$ $\text{unop} ::= \text{not}$
 $v ::= \text{nil} \mid \text{bool_literal}$

Semantics

$$\frac{v \notin \{\text{nil}, \text{false}\}}{\text{if } v \text{ then } s_1 \text{ else } s_2 \xrightarrow{s/e} s_1} \text{ (IF-T)} \qquad \frac{v \in \{\text{nil}, \text{false}\}}{\text{if } v \text{ then } s_1 \text{ else } s_2 \xrightarrow{s/e} s_2} \text{ (IF-F)}$$

$$\text{not } v \xrightarrow{s/e} \delta(\text{not}, v) \text{ (NOT)} \qquad v \text{ binop } e \xrightarrow{s/e} \delta(\text{binop}, v, e) \text{ (BINOP)}$$

$$\delta(\text{and}, v, e) = \begin{cases} v & \text{if } v \in \{\text{nil}, \text{false}\} \\ e & \text{otherwise} \end{cases}$$

$$\delta(\text{or}, v, e) = \begin{cases} v & \text{if } v \notin \{\text{nil}, \text{false}\} \\ e & \text{otherwise} \end{cases}$$

$$\delta(\text{not}, v) = \begin{cases} \text{true} & \text{if } v \in \{\text{nil}, \text{false}\} \\ \text{false} & \text{otherwise} \end{cases}$$

Fig. 1 Pure fragment of Lua₀

numbers, and we will still take this rule as is. Indeed, in Lua, a value like the number 0 is considered **true**, unlike in languages like C.

As for the semantics of expressions, instead of providing operational rules for each operator, we follow [8] and use a function, called *the interpretation function* δ , which provides meaning to operators using a *declarative* style of semantics. That is, we put the emphases on what is the result obtained from using an operator, rather than describing how each operator produces a result step by step. The benefits of using this style are more clearly seen in coming sections when providing semantics to more complex operations than arithmetic operators. To maintain cohesion in our semantics, we will use δ to provide semantics to every primitive operator and service.

In this section, δ provides meaning to the shortcut boolean operators, in which the right operand of a binary operator is not expected to be a value. As with the conditional, a value that is distinct from **false** and **nil** is considered **true**.

A final note regarding the model introduced so far: to present a familiar grammar to the Lua programmer, we maintain the syntactic distinction between statements and expressions. However, we do not let it to scale up to the level of the semantics relations. First, this distinction is merely a concrete grammar issue [23]. Second, previous experience with the formulation of a formal semantics for Lua [37] shows that defining semantics relations following that distinction does not result in an improved understanding of the concepts, but it does result in an unnecessary proliferation of relations.

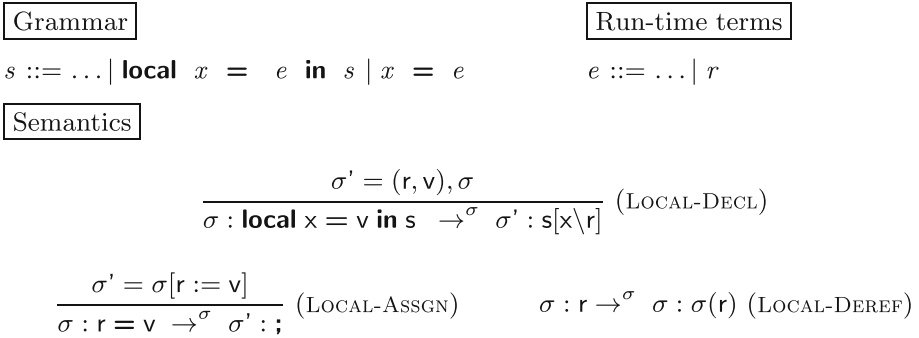


Fig. 2 Stateful fragment of Lua₀

3.2 The Stateful Fragment

We extend the language presented so far with imperative features, namely imperative variables (Fig. 2). Statements are extended with local variable definitions and assignment. There are two things to note about a variable declaration: first, as mentioned in Sect. 2.2, they must be preceded with **local**; otherwise, in Lua, they are global. Second, we make here the first of a few modifications to Lua’s syntax: we explicitly declare the scope of a variable with the **in** keyword, absent in real Lua. This minor change in the grammar allows us to have a very simple representation of state that, as will become clear, brings several benefits.

The operational semantics of imperative features are commonly understood in terms of state and state change; therefore, we enrich our semantics with a model of state: a partial function from a set of references to values, denoted as σ . We refer to σ as the “values’ store,” or simply *store*.

As for the domain of σ , referred as to $dom(\sigma)$, we do not force any specific representation and just require it to be a finite set, with elements that must be syntactically represented and different from any other syntactic object in the language. We further assume it is always possible to obtain a fresh reference from the store. We will write $(r, v), \sigma$ to mean the extension of store σ with reference r pointing to value v , and we will assume that r is *fresh*, i.e., not in the domain of σ .

For the semantics of the new constructs, we use a new relation \rightarrow^σ , which maps a pair of a store σ and a term t (either a statement or an expression) to another pair of a new store σ' and the resulting term t' . The first rule (LOCAL- DECL) models the declaration of a local variable. When the definition is a value v , we put it in the store mapped with a fresh reference r . Then, we replace each occurrence of variable x in the program s with r , denoted as $s[x \setminus r]$.

By substituting every occurrence of the variable with a fresh reference (e.g., [4]), we obtain two benefits: we avoid having to carry around the *environment* (a mapping between variables and references), and we obtain a simple representation for closures (c.f., Sect. 4.2.3). As a little downside, it forces us to add a semantic component to the language, making references be part of the grammar of expressions. We call a *run-time term* such an extension to the language needed to express its semantics. In more advanced sections, they will also serve us to cleanly modularize the semantics.

As for variable assignment (LOCAL- ASSGN), we update the store using the notation $\sigma[r := v]$ to mean a new store in which the value of r is replaced with v . This is only defined if $r \in dom(\sigma)$, which is an implicit side condition of LOCAL- ASSGN that must hold in order to

Evaluation contexts

$$E ::= \llbracket \ \ \rrbracket \mid \text{if } E \text{ then } s \text{ else } s \mid \text{local } x = E \text{ in } s \\ \mid x = E \mid E \text{ binop } e \mid \text{unop } E$$

Standard relation

$$\frac{t \xrightarrow{s/e} t'}{\sigma : E \llbracket t \rrbracket \mapsto \sigma : E \llbracket t' \rrbracket} \text{ (FWD-PURE)} \qquad \frac{\sigma : t \xrightarrow{\sigma} \sigma' : t'}{\sigma : E \llbracket t \rrbracket \mapsto \sigma' : E \llbracket t' \rrbracket} \text{ (FWD-}\sigma\text{)}$$

Fig. 3 Semantics of Lua₀ programs

be able to apply it. Formally, $\sigma[r := v]$ denotes a new store σ' , such that $dom(\sigma') = dom(\sigma)$, where $\sigma'(r) = v$ and $\forall r' \in dom(\sigma'), r' \neq r \Rightarrow \sigma'(r') = \sigma(r')$. Note that the assignment reduces to a skip, indicating that there is nothing else to do for this particular statement. Finally, rule LOCAL- Deref shows that references are implicitly dereferenced.

3.3 Executing Entire Programs

We have already defined two different relations, each of them computing a bit of a program: $\xrightarrow{s/e}$ relates stateless statements or expressions and $\xrightarrow{\sigma}$ stateful ones. Now we are ready to combine the two relations to perform the execution of a complete program. To that effect, we define the \mapsto relation. This relation selects the next term to be executed within a program, called the *redex* and forwards it to the corresponding relation. We are aiming at a deterministic semantics, so there should be only one redex in every program.

We capture the notion of redex syntactically with an *evaluation context* (Fig. 3): a program with a *hole* $\llbracket \ \ \rrbracket$ specifying where the redex is to be found. For every constructor of a term from our language, and for a given position into the term where the next redex must be found, there will be one evaluation context's constructor labeling that position with $\llbracket \ \ \rrbracket$. The determinism of the semantics emanates then from the fact that, for every program, either it is a *final* computation, or there is only one way of decomposing it into an evaluation context and a term, and that term can be executed using only one rule (see Sect. 5).

We can observe from the definition of E the expected evaluation order: in an **if** statement, the guard must be evaluated first. For defining local variables, we evaluate the rvalues first. In a binary operation, we evaluate the left operand first.³ Note that, unlike with the grammar presented so far, for simplicity, we are using a single category of evaluation contexts for statements and expressions.

As stated above, the semantics of programs is given with the \mapsto relation. $E\llbracket t \rrbracket$ denotes an evaluation context where its hole has been filled with the term t , yielding a program. This relation simply forwards execution to the relations defined in previous sections.

3.4 Example

Figure 4 presents the reduction of a simple program that flips boolean variable b . In the right-most column, we write the corresponding evaluation context considered for the step, leaving

³ Our definition enforces left-to-right evaluation of expressions. Even if this is left unspecified in Lua's reference manual, that is how expressions are evaluated in the two most popular implementations of Lua, the reference interpreter and LuaJIT (luajit.org).

σ	t	E
\emptyset	local $b = \text{false}$ in	$[[\]]$
$\mapsto (r1, \text{false})$	if not b then $b = \text{true}$ else $b = \text{false}$	if not $[[\]]$ then ... else ...
$\mapsto (r1, \text{false})$	if not false then $r1 = \text{true}$ else $r1 = \text{false}$	if $[[\]]$ then ... else ...
$\mapsto (r1, \text{false})$	if true then $r1 = \text{true}$ else $r1 = \text{false}$	$[[\]]$
$\mapsto (r1, \text{false})$	$r1 = \text{true}$	$[[\]]$
$\mapsto (r1, \text{true})$	$;$	$[[\]]$

Fig. 4 Trace of execution of a simple program

to the reader the other two components of a reduction step: the fragment of the program filling the hole and the rule being used. Note how in the first step, the variable b is replaced with reference $r1$.

With all of the main ingredients in place, we are now ready to provide semantics to Lua.

4 A Formal Description of Lua

In this section, we describe the highlights of our formalization of the semantics of Lua, the main contribution of this work. Section 4.1 covers the stateless subset of the language, Sect. 4.2 covers the imperative subset, Sect. 4.3 describes the concepts added to support standard library services, Sect. 4.4 covers the semantics of metatables, and Sect. 4.5 completes the semantics with execution order, to be able to explain the execution of complete programs and error handling.

4.1 Stateless Lua

We extend the stateless subset of Lua_0 presented in Sect. 3 with new statements (while loops, breaks, composition of statements) and new values and expressions (numbers and strings, with their corresponding operations). To ease the presentation, we present each addition separately, starting with the new statements.

4.1.1 Statements

Figure 5 shows the new statements. As mentioned, these are compositions of statements (essentially, one statement after the other) and **while** loops (with their **break** operation). We purposely do not include **for** loops and instead translate them into **while** loops, following Lua’s reference manual definition (3.3.5 [22]).

As for the semantics, we start with the composition of statements (SEQ): once the left-most statement has been executed to an end (\dagger), the execution continues with what follows (s).

More interesting is the semantics of while loops. First, a while loop is labeled with a **BREAK** label (**WHILE-START**), renaming **while** to **\$iter**. The labeling and **\$iter** are new run-time statement created for the purpose of executing whiles: the label marks the point from which a **break** should continue the execution, and **\$iter** is necessary to avoid repeatedly unfolding a while, stacking up labels.

A loop marked with **\$iter** is then unfolded as usual, using an **if-then-else** to check the guard and perform a new iteration (**WHILE-ITER**). We took from [8] the idea of labeling the while loop to mark the exit point of a **break**; however, our label is inserted at run time,

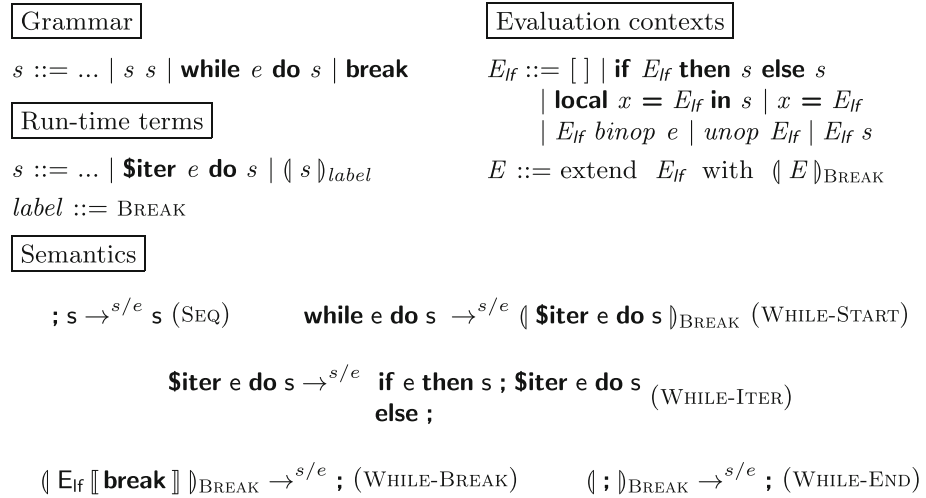


Fig. 5 The stateless subset of Lua (statements)

while in the aforementioned work, it is inserted when *desugaring* the code. Our take reduces the complexity of the desugared code, but at the expense of requiring the extra run-time construct **\$iter** and the labeling step. While in this particular case, the gain is little, in the case of function calls (Sect. 4.2.3), the gain is more evident. For consistency, thus, we decided to do the same for while loops.

When the execution finds a **break** inside a labeled block (**WHILE-BREAK**), the whole program contained in the inner-most **BREAK** label is discarded, effectively modeling the jump out of the body of the loop. To model the program within the label, we use a new category of evaluation context, $E_{\mathbf{if}}$ (*label-free* contexts), which represents a program in which no other labeled term or **while** loop occurs. In order to avoid repetition, we decided to split the evaluation context in two: first, $E_{\mathbf{if}}$ is like the evaluation context from Sect. 3.3, also adding the case for composition of statements; second, we define the new E context as an extension of $E_{\mathbf{if}}$. Extending an evaluation context like $E_{\mathbf{if}}$ to create a new evaluation context E has the effect of renaming each occurrence of $E_{\mathbf{if}}$ with E .

When the execution of the loop ends normally (**WHILE-END**), we just discard the label.

The presented features cover the main aspects of the semantics and control flow of the **while** statement. Note that Lua 5.2 does not have a **continue** statement, yet some of the proposed ways of emulating it⁴ involve concepts already present into our model: using conditionals, Lua’s error mechanism (to be introduced in Sect. 4.5), or using a **break**.

4.1.2 Expressions

We turn our attention to the evaluation of expressions, which we break in two: evaluation of equality and the application of the remaining binary operators. We omit the evaluation of boolean expressions and of the unary minus operator since they follow a similar pattern.

Binary operators, except equality, are treated in Fig. 6. As for the grammar, we define the class `strictbinop` to differentiate binary operators that requires the two values to be evaluated

⁴ See <http://lua-users.org/wiki/ContinueProposal>

Grammar

$v ::= \dots \mid \textit{number_literal} \mid \textit{string_literal}$
 $e ::= \dots \mid e \textit{strictbinop} e$
 $\textit{strictbinop} ::= + \mid - \mid * \mid / \mid ^ \mid \% \mid ..$
 $\quad \mid < \mid \leq \mid > \mid \geq \mid ==$
 $\textit{unop} ::= \dots \mid - \mid \#$

Run-time terms

$e ::= \dots \mid (e)_{\textit{label}}$
 $\textit{label} ::= \dots \mid \text{BINOPWO} \mid \text{EQFAIL}$

Evaluation contexts

$E_{\textit{lf}} ::= \dots \mid E_{\textit{lf}} \textit{strictbinop} e$
 $\quad \mid v \textit{strictbinop} E_{\textit{lf}}$

Semantics of binary operators

$$\frac{\text{op} \in \{+, -, *, /, ^, \%, \dots, <, \leq\} \quad \delta(\text{op}, v_1, v_2) \in v}{v_1 \text{ op } v_2 \rightarrow^{s/e} \delta(\text{op}, v_1, v_2)} \text{ (BINOP)}$$

$$\frac{\text{op} \in \{+, -, *, /, ^, \%, \dots, <, \leq\} \quad \delta(\text{op}, v_1, v_2) \notin v}{v_1 \text{ op } v_2 \rightarrow^{s/e} (v_1 \text{ op } v_2)_{\text{BINOPWO}}} \text{ (BINOP-WO)}$$

Fig. 6 The stateless subset of Lua (binary operators)

prior to execution, unlike with the boolean operators. Note how we extend the evaluation context $E_{\textit{lf}}$, and assume that E is extended likewise.

Turning to their semantics, BINOP states that the actual semantics of the application of op over operands v_1 and v_2 is captured by the function $\delta(\text{op})$. The semantics of the binary operators shown is mostly standard, the only exception being an internal step of coercion done between numbers and strings: in every Lua’s primitive operator and library service, each time that the semantics dictates that a number is expected, it is possible to provide a string representation of a number, and internally, Lua tries to coerce it to a number, following certain lexical conventions. The same happens when a string is expected and a number is provided. These details, together with the remaining aspects of the standard semantics of binary operations, are all captured in δ . The only role of the semantics relations is to interface with δ , checking that it is actually possible to apply op over operands v_1 and v_2 , *i.e.*, the application returns a value v , and not, for example, an error object.

On the other hand, rule BINOP-WO shows the case when the binary operator cannot be applied successfully ($\delta(\text{op}, v_1, v_2) \notin v$). This could mean that we tried to perform arithmetic over values that cannot be interpreted as numbers, or tried to compare values for which there is no standard order relation built-in in Lua. In this case, we label the expression with information about what has gone wrong (BINOPWO, where WO stands for Wrong Operands). We, therefore, extend the class of expressions to include labeled expressions and extend the set of labels with the new labels, in Fig. 6.

At this point, execution is stuck here. If we were not to extend the semantics any further, these error labels will only serve the purpose of informing the user what caused her program to fail. Later on, we will see how metatables try to resolve the situation, either by providing an alternative value, or by *throwing* an error (Sect. 4.4). Such errors can be caught with the primitive `xpcall` (Sect. 4.5).

In case of equality, the behavior is a bit more intricate (Fig. 7). As with previous operators, we abstract into $\delta(==)$ the corresponding details of equality comparison, but we perform some extra processing of the result. However, note that the equality relation is not trivial: *e.g.*, in

$$\begin{array}{c}
 \frac{\delta(==, v_1, v_2) = \mathbf{true}}{v_1 == v_2 \xrightarrow{s/e} \mathbf{true}} \quad (\text{EQ-TRUE}) \\
 \\
 \frac{\delta(==, v_1, v_2) = \mathbf{false} \quad \delta(\text{type}, v_1) \neq \text{"table"} \vee \delta(\text{type}, v_2) \neq \text{"table"}}{v_1 == v_2 \xrightarrow{s/e} \mathbf{false}} \quad (\text{EQ-FALSE}) \\
 \\
 \frac{\delta(==, v_1, v_2) = \mathbf{false} \quad \delta(\text{type}, v_1) = \text{"table"} \quad \delta(\text{type}, v_2) = \text{"table"}}{v_1 == v_2 \xrightarrow{s/e} (\downarrow v_1 == v_2)_{\text{EQFAIL}}} \quad (\text{EQ-FAIL})
 \end{array}$$

Fig. 7 The stateless subset of Lua (equality)

order to obtain a semantics that can be tested against the official test suites, it must define proper comparison of IEEE 754 floating point numbers (as in a standard binary of the official interpreter).

In essence, $\delta(==)$ compares numbers and strings by value, and tables and closures by reference—with values of different types being always distinct. That is, unlike other dynamic languages like JavaScript, $0 == "0"$ evaluates to **false**.

When the objects are equal (EQ- TRUE), we simply return **true**. But if the objects are distinct (for instance, they have distinct types), we have two options: if the operands involved are not tables (EQ- FALSE), equality comparison reduces to **false**. Otherwise, if they are tables (EQ- FAIL), instead of returning **false**, we label the expression and defer its result to the metatable mechanism (Sect. 4.4). The metatable mechanism will then look up for user-defined functions (*metamethods*) implementing a comparison method for non-equal table values. In that way, Lua allows for the definition of arbitrary (reflexive) relations over table values.

4.2 Imperative Lua

In this section, we grow the imperative fragment presented in Sect. 3 to include tables (Sect. 4.2.1), multiple local variables definition and assignment (Sect. 4.2.2), and functions (Sect. 4.2.3). The latter is presented here for reasons that will become clear in time.

4.2.1 Tables

The syntax of table constructors, table field indexing and assignment, is presented in Fig. 8, together with their associated semantic rules. Essentially, they are mutable associative arrays: data structures where we can store key/value pairs, where keys cannot be **nil** or **nan** (a value in the set of numbers used for undefined or not representable numeric results), and values cannot be **nil**.

Tables are manipulated through references, so, to that end, we create a new mapping between tables' identities (*tid*) and tables (and associated meta-data), called θ . The elements from the domain of θ are considered values and satisfy the same properties as references in σ , with the reasonable addition that the former should be syntactically distinguishable from the latter. The image of θ only contains tables and any table-related meta-data. For the present model, this meta-data reduce to information about the associated metatable, introduced later. Formally, $\theta : tid \rightarrow tables \times \{\mathbf{nil}, tid\}$, where *tables* denotes the set of tables as described by the grammar in Fig. 8, and **nil** or *tid* are the possible metatable-related values associated

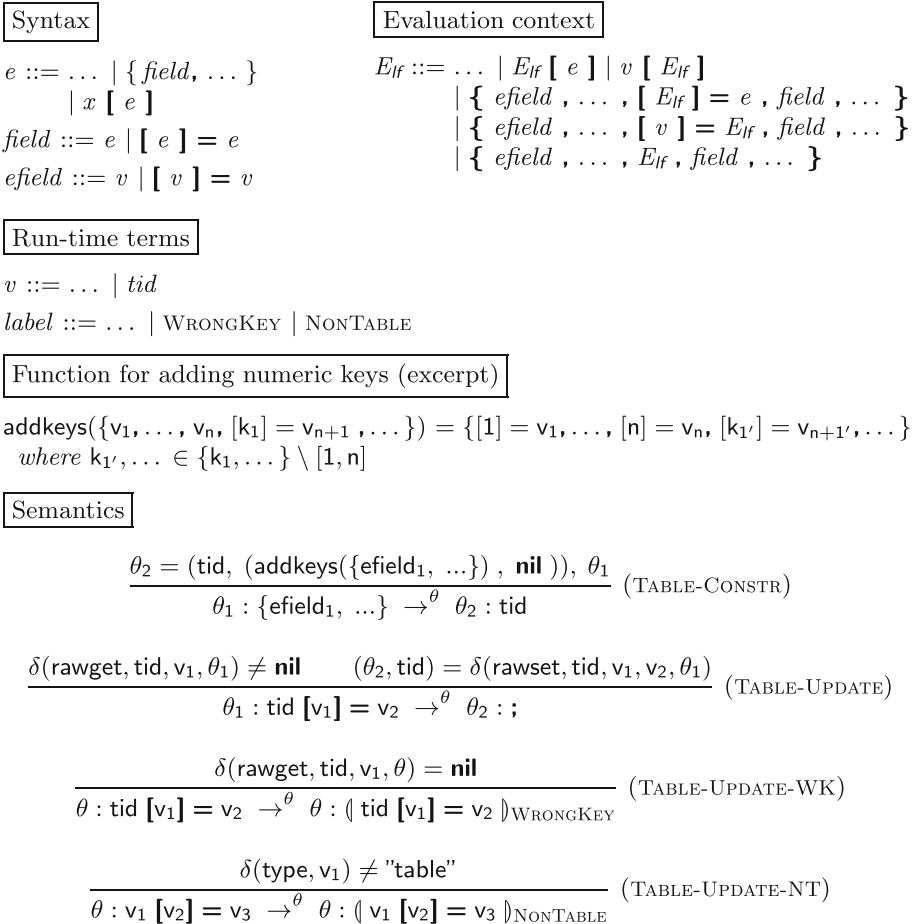


Fig. 8 Formalization of tables (excerpt)

with the given table, to be introduced later. We capture a new store θ resulting from, either, extending a previous store with a new pair tid /table or changing the content of a previously stored table, using the same devices and notation as with σ (Sect. 3.2).

The construction of a table is formalized in Rule TABLE- CONSTR (Fig. 8), which acts when all of its fields have been evaluated. Fields are evaluated from left to right, with each key being evaluated before each value. Note the new syntax non-terminal *efield* to denote evaluated fields, which can be just a value or a key/value pair (denoted $[v] = v$). The meta-function *addkeys* adds absent keys—namely, consecutive natural numbers—to fields of the form v . It also discards fields already present in the constructor that contain numeric keys in the interval $[1, n]$. Here, n is the quantity of fields without keys in the constructor (hence, the fields for which *addkeys* adds the corresponding numeric keys). We show an excerpt of its formal definition, in Fig. 8. Note that Lua’s reference manual is not precise in this point, and different definitions of *addkeys* can have an observable impact on the semantics of operations like computing the *length* of a table (through the operator #), or iterating a table,

$$\delta(\text{rawget}, v_1, v_2, \theta) = \delta(\text{error}, \dots) \quad \text{if } \delta(\text{type}, v_1) \neq \text{"table"}$$

$$\delta(\text{rawget}, \text{tid}, v_i, \theta) = \begin{cases} v_j & \text{if } \theta(\text{tid}) = (\{\dots, [v'_i]=v_j, \dots\}, \dots) \\ & \text{and } \delta(==, v_i, v'_i) = \text{true} \\ \mathbf{nil} & \text{otherwise} \end{cases}$$

$$\delta(\text{rawset}, \text{tid}, v_1, v_2, \theta_1) = (\theta_2, \text{tid}),$$

$$\text{where } \begin{cases} v_1 \notin \{\mathbf{nil}, \mathbf{nan}\} \\ \theta_1(\text{tid}) = (\{\dots, \text{efield}_{i-1}, [v'_1]=v_3, \text{efield}_{i+1}, \dots\}, \dots) \\ \delta(==, v_1, v'_1) = \text{true} \\ \theta_2 = \begin{cases} \theta_1[\text{tid} := (\{\dots, \text{efield}_{i-1}, \text{efield}_{i+1}, \dots\}, \dots)] & \text{if } v_2 = \mathbf{nil} \\ \theta_1[\text{tid} := (\{\dots, \text{efield}_{i-1}, [v'_1]=v_2, \text{efield}_{i+1}, \dots\}, \dots)] & \text{if } v_2 \neq \mathbf{nil} \end{cases} \end{cases}$$

$$\delta(\text{rawset}, \text{tid}, v, v', \theta_1) = (\theta_2, \text{tid}),$$

$$\text{where } \begin{cases} v \notin \{\mathbf{nil}, \mathbf{nan}\} \\ \theta_1(\text{tid}) = (\{[v_1]=v'_1, \dots, [v_n]=v'_n\}, \dots) \\ \forall k \in \{v_1, \dots, v_n\}, \delta(==, v, k) = \text{false} \\ \theta_2 = \begin{cases} \theta_1 & \text{if } v' = \mathbf{nil} \\ \theta_1[\text{tid} := (\{[v]=v', [v_1]=v'_1, \dots, [v_n]=v'_n\}, \dots)] & \text{if } v' \neq \mathbf{nil} \end{cases} \end{cases}$$

$$\delta(\text{rawset}, v_1, v_2, v_3, \theta) = (\theta, \delta(\text{error}, \dots)), \quad \text{if } \delta(\text{type}, v_1) \neq \text{"table"} \vee v_2 \in \{\mathbf{nil}, \mathbf{nan}\}$$

Fig. 9 Formalization of tables (primitives for table assignment and indexing)

through library services like `ipairs`. Our mechanization of `addkeys` successfully passes the official interpreter’s tests, and as such, can serve as a frame for future implementations.

Finally, returning to `TABLE- CONSTR` (Fig. 8), note that we are actually mapping a fresh table id `tid` with an ordered pair: the first component contains the actual table and the second one the id of the metatable or `nil`, if absent, as is in this case. Metatables are discussed in depth in Sect. 4.4.

Assignment to table fields has three cases: the successful update of the field (`TABLE- UPDATE`), and two special scenarios. The first comes from trying to update a non-existing key (`TABLE- UPDATE- WK`); while the second comes from trying to update an object which is not a table (`TABLE- UPDATE- NT`). The semantics of these special scenarios can be meta-programmed, so, as we did with similar situations, we just label the term with information about what has occurred, to delegate the execution to the metatables mechanism.

To know if a key exists in a table and to update the field, we use two services of Lua, `rawget` and `rawset`, formalized in Fig. 9. $\delta(\text{rawget}, \text{tid}, v, \theta)$ yields either the value associated with `v` in $\theta(\text{tid})$ or `nil`, if there is no associated value. Note that we use the equality comparison function ($\delta(==)$) to determine which field is being indexed with the given key. Recall from Sect. 4.1.2 that $\delta(==)$ should define the details of a proper comparison between Lua values, and, also, should not resort to the metatables mechanism when comparing different table values (which is not the case with the language `==` comparison operator).

$\delta(\text{rawset}, \text{tid}, v_1, v_2, \theta)$ yields a new θ where the table referenced by `tid` associates `v2` with value `v1`. In the first equation, we consider an existing key, in which case we remove or update the field, depending of the value `v2`. Note how `nil` and `nan` cannot be the keys of a table field. As with `rawget`, the keys are compared using $\delta(==)$, and the same considerations apply with

Grammar

$s ::= \dots \mid \text{var} , \dots = e , \dots$
 $\mid \text{local } x , \dots = e , \dots \text{ in } s$
 $\text{var} ::= x \mid e [e]$
 $\text{evar} ::= r \mid v [v]$

Evaluation context

$E_{lf} ::= \dots$
 $\mid \text{local } x , \dots = v , \dots , E_{lf} , e , \dots \text{ in } s$
 $\mid \text{evar} , \dots , E_{lf} [e] , \dots = e , \dots$
 $\mid \text{evar} , \dots , v [E_{lf}] , \dots = e , \dots$
 $\mid \text{evar} , \dots = v , \dots , E_{lf} , \dots$

Semantics

$$\frac{k \geq 1 \quad \text{v}_{n-k+1} = \mathbf{nil}, \dots, \text{v}_n = \mathbf{nil}}{\text{evar}_1, \dots, \text{evar}_n = \text{v}_1, \dots, \text{v}_{n-k} \rightarrow^{s/e} \text{evar}_1, \dots, \text{evar}_n = \text{v}_1, \dots, \text{v}_{n-k}, \text{v}_{n-k+1}, \dots, \text{v}_n} \text{ (ASSGN-FEWER)}$$

$$\frac{k \geq 1}{\text{evar}_1, \dots, \text{evar}_n = \text{v}_1, \dots, \text{v}_n, \dots, \text{v}_{n+k} \rightarrow^{s/e} \text{evar}_1, \dots, \text{evar}_n = \text{v}_1, \dots, \text{v}_n} \text{ (ASSGN-MORE)}$$

$$\frac{n \geq 2}{\text{evar}_1, \dots, \text{evar}_n = \text{v}_1, \dots, \text{v}_n \rightarrow^{s/e} \text{evar}_1 = \text{evar}_1 \quad \text{evar}_2, \dots, \text{evar}_n = \text{v}_2, \dots, \text{v}_n} \text{ (ASSGN-SPLIT)}$$

$$\frac{\sigma' = (r_1, \text{v}_1), \dots, (r_n, \text{v}_n), \sigma}{\sigma : \text{local } x_1, \dots, x_n = \text{v}_1, \dots, \text{v}_n \text{ in } s \rightarrow^\sigma \sigma' : s[x_1 \setminus r_1, \dots, x_n \setminus r_n]} \text{ (LOCAL-DECL)}$$

Fig. 10 Multiple variables declaration and assignment (excerpt)

regard to the semantics consequences when using user-defined equality comparison relations. The second equation of `rawset` considers the case when the key does not exist in the table. In this case, the pair is added, only if the value is not `nil`. The last equation considers the error cases.

Accessing a field of a table follows a similar pattern so we omit the rules of this operation.

4.2.2 Local Variables

We introduce a small change in the language in order to support multiple-variable definition and assignment (Fig. 10). Essentially, we can have arbitrary lists of variables and terms as the lvalues and rvalues of a local declaration, respectively. Similarly, we can have arbitrary lists in the assignment, with the addition that we can update variables *and* tables' fields.

For multiple-variable assignment, the following steps are performed. First, the execution evaluates all lvalues before the rvalues, as indicated by the evaluation contexts. Second, the quantity of rvalues is adapted to match the lvalues, either adding extra `nil` rvalues when there are fewer rvalues than lvalues (ASSGN- FEWER), or discarding the extra rvalues when there are more rvalues than lvalues (ASSGN- MORE). Once the number of lvalues matches that of rvalues, a multiple assignment is decomposed into a sequence of single assignments. This is so since rvalues can be local variables and/or table fields, and we need to distinguish each case in order to properly explain their semantics. Note that we need not to make any further step: each assignment will be handled by the corresponding rules (LOCAL- ASSGN, from Sect. 3.2, and TABLE- UPDATE, from Sect. 4.2.1). Finally, by splitting a multiple assignment into several single variable assignments, we also model what happens in Lua when there is an lvalue that is repeated: the assignments are effectively repeated over the local variable or table field.

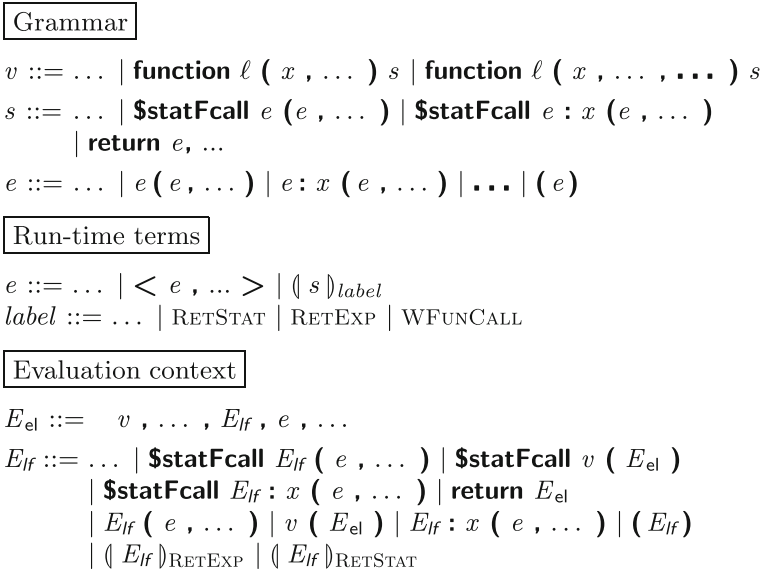


Fig. 11 Functions and function calls (grammar)

The introduction of local variables follows the same steps, which we omit for brevity. The creation of new references is modeled with a simultaneous substitution operation (LOCAL-DECL, which subsumes the previously defined rule with the same name in Sect. 3.2). Repeated occurrences of variable identifiers (as lvalues) are allowed, and their correspondent values are evaluated. Then, when the introduction of the variable in the corresponding scope is applied (through the substitution function), only the last variable definition is taken into account.

4.2.3 Functions, Function Applications, and Tuples

It might at first look suspicious to consider Lua’s functions and application as belonging to the imperative subset. However, Lua’s functions are first-class values internally represented with closures closing over their external variables’ references, and, also, the parameters of a function are modeled as imperative variables. Therefore, we need the state to define functions and function’s applications. Tuples are also considered here because they are required to model functions.

Figure 11 shows the grammar involved in the formalization of function definitions and calls. The first thing to note is that functions are values, and that each function is labeled with a new countable set of elements ℓ , which should not be confused with a function’s name: the label ℓ —inserted at *desugaring*—is used for closure comparison (Sect. 4.2.4), and in practice, it is akin to the line number, in the source code. As an example, **function** $f ()$ **end** is desugared as $f = \mathbf{function} \ell_1 () \mathbf{end}$.

Functions have two constructors; the difference being the *vararg* arguments \dots (we use this special notation to avoid confusion with ellipses also used in the grammar; when showing Lua code, we will use the conventional syntax \dots). This special token refers to the extra arguments passed to a function beyond of those explicitly named in its definition. Within the

$$\begin{array}{c}
 \frac{1 \leq i \leq \min(m, n) \Rightarrow v'_i = v_i \quad i > m \Rightarrow v'_i = \mathbf{nil} \quad \sigma' = (r_1, v'_1), \dots, (r_n, v'_n), \sigma}{\sigma : (\mathbf{function} \ell (x_1, \dots, x_n) s) (v_1, \dots, v_m) \rightarrow^{\sigma} \quad \sigma' : \llbracket s [x_1 \setminus r_1, \dots, x_n \setminus r_n] \rrbracket_{\text{RETEXP}}} \quad (\text{E-CALL}) \\
 \\
 \frac{1 \leq i \leq \min(m, n) \Rightarrow v'_i = v_i \quad i > m \Rightarrow v'_i = \mathbf{nil} \quad \text{tuple} = \langle v_{n+1}, \dots, v_m \rangle \quad \sigma' = (r_1, v'_1), \dots, (r_n, v'_n), \sigma}{\sigma : (\mathbf{function} \ell (x_1, \dots, x_n, \dots) s) (v_1, \dots, v_m) \rightarrow^{\sigma} \quad \sigma' : \llbracket s [x_1 \setminus r_1, \dots, x_n \setminus r_n, \dots \setminus \text{tuple}] \rrbracket_{\text{RETEXP}}} \quad (\text{E-CALLVARG}) \\
 \\
 \frac{\delta(\text{type}, v) \neq \text{"function"}}{v (v_1, \dots, v_n) \rightarrow^{s/e} \quad \llbracket v (v_1, \dots, v_n) \rrbracket_{\text{WFUNCALL}}} \quad (\text{E-CALLWRONG}) \\
 \\
 v:\text{name} (e_1, \dots, e_n) \rightarrow^{s/e} v[\text{"name"}] (v, e_1, \dots, e_n) \quad (\text{E-MCALL}) \\
 \\
 \llbracket E_{\text{if}} \llbracket \mathbf{return} (\mathbf{function} \ell (x_1, \dots) s) (v_1, \dots) \rrbracket_{\text{RETEXP}} \quad (\mathbf{function} \ell (x_1, \dots) s) (v_1, \dots) \rrbracket_{\text{RETEXP}} \rightarrow^{s/e} \quad (\text{E-POPSF})
 \end{array}$$

Fig. 12 Functions and function calls (semantics, 2nd step)

function body, ... is treated as a *tuple*. Tuples are run-time expressions, noted as $\langle e, \dots \rangle$, and cannot be constructed by the user. They are not values and, hence, cannot be stored.

Then, we have the **return** statement, which allows the return of multiple values; function calls, with a special syntax for calling an object’s method; and *parenthesized expressions* (e), which takes the first value of a list of values.

As for function calls, we need to distinguish their occurrences: they can appear in a statement, as in `f()`; `print('hi')`, or in an expression, as in `print(f())`. A function call used as an expression can return none, one or more values, while a function call used as a statement is useful just for its side effects, with any value returned from such function being discarded. Our parser is in charge of distinguishing and compiling each case into syntactically different objects, tagging with **\$statFcall** the latter case.⁵

Function calls are not plain β -contractions: they allow us to call a given function with an arbitrary quantity of parameters, regardless of what is specified in its signature, discarding or completing with extra **nil** values as required, like the following simple example shows:

Example 8 (Calling a function with different number of parameters.)

```

function f(x1, x2) return x1, x2 end
print(f())      -->> nil  nil
print(f(1))     -->> 1   nil
print(f(1,2))   -->> 1   2
print(f(1,2,3)) -->> 1   2
    
```

Additionally, the result of a call might be in a position where the returned values must be appended to existing ones or trimmed. For instance, consider the following short example, in which maximum is the function from Example 5:

Example 9 (Returned values are appended to the argument list or trimmed.)

```

m = maximum({1, maximum{1,2,2,3}})    -- m: 4
    
```

⁵ This distinction helped simplify the model, as we do not need to define new notions of evaluation contexts to distinguish among either case—a task that proved to be cumbersome and difficult to maintain through the evolution of the model.

Remember that `maximum` returns the greatest number of a list together with its position. To understand this result, first note that the inner call returns 3 and 4. These values are then appended to the outer list, resulting in {1, 3, 4}. Therefore, the second call to `maximum` yields values 4 and 3, the second value being trimmed out and the first one being assigned to `m`.

A last aspect of function calls that can be captured into the model presented so far, is related with *tail calls*: nested function calls can reuse the stack frame of the actual called function, if the function call is the last instruction.⁶ Specifically, in Lua 5.2, tail calls are calls of the form **return** (function ℓ (x_1, \dots) s) (v_1, \dots). In that case, the called function will reuse the stack frame of the actual function from which it is being called. In our model, we include just one piece of information from a regular stack frame: the position to which the function must return, denoted with the labels `RETSTAT` and `RETEXP`. And by means of the manipulation of this representation, we can model tail calls. We explain it below.

In order to specify these behaviors, we take the following actions:

1. Evaluate the function position and then the arguments, from left to right.
2. Recognize function calls in tail position and act accordingly.
3. Adjust the list of actual parameters and perform the function call itself.
4. According to the place of the call (statement or expression), discard the returned values or put them in a tuple.
5. In the case of an expression, distinguish if the tuple must be appended to other values or trimmed.

The first step is specified with the help of the evaluation contexts, as can be seen in Fig. 11. As a shorthand, we introduce a new non-terminal E_{el} to mean a list of arguments whose elements to the left of the evaluation point are fully evaluated.

The second step is partially described by in Fig. 12, last rule, only for the case of a nested function call made from another function call embedded into an expression. In order to better understand it, we first explain the next step.

The third step is shown in Fig. 12, only for the case of function calls in expressions, as the other case is analogous. Rule E-CALL models a simple call, with the parameters of the function treated as mutable variables, so a fresh reference is used to model each parameter. If there are fewer arguments, the value `nil` is assigned to the remaining parameters of the function, and if there are more, they are simply silently ignored. The resulting term—the body of the function with the fresh references replacing the parameters—is labeled with a new run-time expression $(\cdot)_{RETEXP}$. This labeling serves two purposes: first, it allows to recognize the function call as, in this case, an expression, something needed when reasoning about tuples (see below); and second, it marks the place where a **return** statement must jump to. In the case of a function call in a statement, the label is `RETSTAT`.

Rule E-CALLVARG shows the case of a vararg function call, with the difference residing on what is done with surplus arguments: instead of being ignored they are put into a tuple, which replaces the vararg expression (\dots) in the body of the function. Note that they are not stored, as it does occur with the remaining actual parameters. Formally, we treat the vararg mark occurrence in the signature of a function as a binding occurrence, with special scoping rules in order to account for a particularity discussed in ([30], Issue #1): the binding occurrence binds any vararg mark that occurs in the body of the function, except if it appears in the body of a nested function definition, regardless of it being or not a vararg function. The following interaction with the official interpreter (taken from [30], Issue #1) illustrates this scoping rule:

⁶ We thank the reviewers for inviting us to enrich our model with this and other features.

```
> function f( ... )
  return function() return ... end
end
```

stdin :1: cannot use ' ... ' outside a vararg **function** near ' ... '

Rule E- CALLWRONG describes the exceptional situation that might be handled by the metatable mechanism: a function calls over a non-function value. As with previous special situations, we just label the whole expression with a tag documenting what happened.

Finally, E- MCALL describes method invocation: it is simply translated into a table lookup, with the object being injected as the first argument of the function.

Now, previous to the function call itself, we try to recognize if it is in tail position. Remember that, in Lua 5.2, tail calls reduce to calls of the form:

```
return (function  $\ell$  ( $x_1, \dots$ )  $s$ ) ( $v_1, \dots$ )
```

This step will be responsibility of the top level relation \mapsto : it will be in charge of guaranteeing that a tail call is correctly processed before actually executing the function call itself. Once that a tail call is recognized, we manipulate the stack of calls: in our model, we just look for the inner-most RETSTAT or RETEXP label around the function call, and remove it together with the whole rest of computation between the label and the tail call itself. That is, we remove the computations that would be discarded when performing the return of the function being called in tail position. More formally, we look for a function call that looks like this, for the case of a tail call nested into another function call done in the position of an expression:

```
(  $E$ if [ return (function  $\ell$  ( $x_1, \dots$ )  $s$ ) ( $v_1, \dots$ ) ] )RETEXP
```

In that case, and in terms closer to an actual implementation, we could say that we *pop* the actual stack frame from the call stack, obtaining

```
(function  $\ell$  ( $x_1, \dots$ )  $s$ ) ( $v_1, \dots$ )
```

This is performed by rule E- POPSF, in Figure 12. Then, the actual execution of the function call, as described by the previous semantics rules, will install a new stack frame for the current tail call. Together, all these steps model the phenomenon of a tail call.

A last remark: in the context of a programming language with interfaces with the garbage collector, like finalizers, these implementation details have an observable effect on the semantics of programs. As an example, consider the following program:

```
function f(a)
  local b = {}
  return f(b)
end

local a = {}
f(a)
```

In the previous program, given that the recursive call to **f** is in tail position, at each recursive call the whole content of the stack frame of the previous call is discarded. This implies that, at each recursive call, the reference to the parameter **a** from the previous call is lost (except for the first call, made in the last line). On the other hand, if the recursive call to **f** would be just **f** (), instead of **return f** (), then the call would not be in tail position, and every stack frame of every call would be preserved, including the reference to the parameter **a**. As we

$$\begin{aligned}
 (\ ; \)_{\text{RETSTAT}} &\rightarrow^{s/e} \ ; \ ; \ (\text{S-RETSKIP}) & (\ E_{\text{if}} [\ \mathbf{return} \ v, \dots \] \)_{\text{RETSTAT}} &\rightarrow^{s/e} \ ; \ ; \ (\text{S-RETURN}) \\
 (\ ; \)_{\text{RETEXP}} &\rightarrow^{s/e} \ < \ > \ (\text{E-RETSKIP}) \\
 (\ E_{\text{if}} [\ \mathbf{return} \ v, \dots \] \)_{\text{RETEXP}} &\rightarrow^{s/e} \ < \ v, \dots \ > \ (\text{E-RETURN}) \\
 (\ E_{\text{if}} [\ \mathbf{return} \ v, \dots \] \)_{\text{BREAK}} &\rightarrow^{s/e} \ \mathbf{return} \ v, \dots \ (\text{S-RETURNBREAK})
 \end{aligned}$$

Fig. 13 Functions and function calls (semantics, 3rd step)

explain in [38] and Sect. 7, Lua 5.2 has a reachability-based garbage collector that, in each recursive tail call of *f*, would consider unreachable the parameter of the previous call. Given that we have interfaces with the garbage collector, this difference in the internal treatment of the stack frames could be observed at the semantics level.

The fourth step is split in two. First, the body of the function is executed to an end, either a skip or a **return**. Second, as listed in Fig. 13, the values are processed as follow: a function call as a statement simply returns skip (S- RETSKIP), discarding the values if needed to (S- RETURN). Instead, in an expression, it returns a tuple, either empty (E- RETSKIP) or containing the values (E- RETURN). A **return** might occur inside the body of a while (S- RETURNBREAK), in which case we must jump outside of the loop.

The fifth step kicks in once we have a tuple. As noted in Example 9, we need to know if values are appended to some existing list of values or trimmed to just obtain the first value. The process requires some machinery but is simple in essence. If the *context*, in which the tuple is immersed, requires a list of values, the tuple’ values are appended to the values existing in the context, else, the first value is taken. To know the context, as the name suggests, we use the same type of contexts used for evaluation.

Figure 14 shows the details of the fourth step. Rule TUPLE- ONE replaces the tuple with its first value. For this rule to be executed, the tuple must be in a context E_t (the *t* stands for *truncate*), which for brevity is shown partially. This means that the tuple is, for instance, in the conditional of an **if-then-else**, or in a list of values (but not at the end), or in a parenthesized expression, *etc.* If the tuple is empty, under the same context, it is replaced with **nil** (TUPLE- ZERO).

Rule TUPLE- APPEND shows the case when the tuple must be appended to a list. This rule uses a new context, E_a (a for *append*), which only applies if the tuple is at the end of a list of values. Then, the result of the tuple is computed with the meta-function **append**, which is also shown in abridged form. This function takes the context and the tuple and simply insert the values of the tuple to the end of the list, or does nothing if the tuple is empty.

With the semantics shown in this section, we can understand the example from the introduction (Example 1). We explain it in two bits. First, we have function *f* returning **nil**:

```

function f() return nil end
print(f(), f(), f()) -->> nil nil nil
    
```

The execution of the **print** statement occurs after each call to *f* is resolved. Since they are in the positions of expressions, after E- CALL, we get $(\ \mathbf{return} \ \mathbf{nil} \)_{\text{RETEXP}}$. With an empty E_{if} context, we perform one E- RETURN step and get a tuple with a **nil** inside. In the first two calls to *f*, the tuple is in a position where a value is expected, therefore after executing TUPLE- ONE we simply get **nil**. In the third call, instead, we append the tuple to the values

Contexts for tuple processing (excerpt)

$$E_{tel} ::= v, \dots, [], e_1, e_2 \dots$$

$$E_t ::= \text{if } [] \text{ then } s \text{ else } s \mid \text{local } x, \dots = E_{tel} \text{ in } s \mid \text{evar}, \dots = E_{tel} \mid ([]) \mid \{ \text{efield}, \dots, [], \text{field}_1, \text{field}_2, \dots \} \mid \dots$$

$$E_{ael} ::= v, \dots, []$$

$$E_a ::= \text{local } x, \dots = E_{ael} \text{ in } s \mid \text{evar}, \dots = E_{ael} \mid \{ \text{efield}, \dots, [] \} \mid \dots$$

Appending function (excerpt)

$$\text{append}(\text{evar}, \dots = v_1, \dots, [], < v_2, v_3, \dots >) = \text{evar}, \dots = v_1, \dots, v_2, v_3, \dots$$

$$\text{append}(\text{evar}, \dots = v_1, \dots, [], <>) = \text{evar}, \dots = v_1, \dots$$

Semantics

$$E_t \llbracket < v_1, v_2, \dots > \rrbracket \rightarrow^{s/e} E_t \llbracket v_1 \rrbracket \text{ (TUPLE-ONE)}$$

$$E_t \llbracket < > \rrbracket \rightarrow^{s/e} E_t \llbracket \text{nil} \rrbracket \text{ (TUPLE-ZERO)}$$

$$E_a \llbracket < v_1, \dots > \rrbracket \rightarrow^{s/e} \text{append}(E_a, < v_1, \dots >) \text{ (TUPLE-APPEND)}$$

Fig. 14 Functions and function calls (semantics, 4th step)

computed so far (TUPLE- APPEND). Therefore, we end with **print(nil, nil, nil)**, which in turn prints the three **nils**.

Different is the case when we do not return anything:

```
function g() end
print(g(), g(), g()) --->> nil nil
```

Note that the body of **g** is an implicit skip. Therefore, in the first two calls to **g** we obtain a skip that is converted to an empty tuple by **E- RETSKIP**, and then to **nil** by **TUPLE- ZERO**. In the last call, as in the previous case, **TUPLEAPPEND** applies and **append** discards the empty tuple. Therefore, we now end with **print(nil, nil)**.

4.2.4 Function Comparison and Caching

Lua 5.2 admits an unspecified caching of closures: when their definitions have *no observable difference* ([13], Chapter 8.1). We implement this caching in a way that is compatible with the official interpreter, correctly mimicking the following interactions:⁷

Example 10 (Caching of functions.)

```
1 f = function() end
2 g = function() end
3 print(f == g) --->> false
4 h = function() return function() end end
```

⁷ Note that LuaJIT 2.0, the other major implementation of (part of) Lua 5.2, does not perform caching of closures.

```
5 print(h() == h()) -->> true
```

The first two functions are equal; however, Lua creates a closure for each of them, as line 3 shows. But when the function is created from the same function definition, it does cache its definition (lines 4–5).

As mentioned at the beginning of Sect. 4.2.3, our function’s definitions include a label ℓ , which is what distinguishes different occurrences of definitions in the program. Observe the partially unsugared code obtained from Example 10:

```
f = (function $1 () ; end)
g = (function $2 () ; end)
print(f == g)
h = (function $3 () return (function $4 () ; end) end)
print(h() == h())
```

Each function created from calling h shares its label ($\$4$), and therefore, we can treat each instance as the same. But this is not enough to identify a function, as the following example shows:

Example 11 [Same function yet different closures.]

```
for i=1,10 do print(function() return i end) end
-->> function: 0x1c4d970 function: 0x1c4e3d0 ...
```

Note how at each iteration it prints a new address: the function is not cached. The reason is that its body refers to a variable, i , which is different at each iteration (**for** loops makes the iteration variable immutable by creating a fresh variable at each iteration). The environment of a closure, including every local variable in which scope reaches to the given closure, is part of what the official interpreter recognizes as “observable difference” (besides the function definition itself), and, accordingly, differences in a scope are what triggers the creation of new closures, from the same function definition. In our formalization, we consider, as a representation of a closure, the whole function definition with its embedded environment, as it is when the focus of reduction reaches to the function definition (see Corollary 2, Sect. 5). Therefore, returning to the previous example, we would include in the corresponding representation of the closures created, each new instance from each iteration of the variable i . At each iteration, we would end up with a different closure and, hence, no closure would be reused.

4.3 Built-in Services

As we saw throughout this work, the interpretation function δ defines the semantics of built-in services (**type**, **rawget**, ...) and operators (**+**, **..**, ...). Originally introduced in ISWIM ([5]) for the purpose of abstracting the semantics over a set of primitive operators, we use it here motivated by different purposes. As mentioned in Sect. 3, instead of describing *operationally* the semantics of built-in services and primitive operators, we use a declarative style; *i.e.*, we specify just the result of the call to a service or the application of an operator, instead of describing how we obtain such results. This approach reflects the style of the reference manual: there is no operational specification from which the results can be obtained, just a description of what is expected from calling a given service or applying a given operator.

Grammar

$svc ::= \text{assert} \mid \text{error} \mid \dots$
 $e ::= \dots \mid \text{\$builtin } svc(e, \dots)$

Semantics

$$\frac{svc \in \{\text{type, assert, error, pcall, select, } \dots\}}{\text{\$builtin } svc(v_1, \dots, v_n) \rightarrow^{s/e} \delta(svc, v_1, \dots, v_n)} \text{ (SVC-NOSTATE)}$$

$$\frac{svc \in \{\text{ipairs, next, pairs, getmetatable, } \dots\}}{\theta : \text{\$builtin } svc(v_1, \dots, v_n) \rightarrow^\theta \theta : \delta(svc, v_1, \dots, v_n, \theta)} \text{ (SVC-CONST)}$$

$$\frac{svc \in \{\text{rawset, setmetatable}\} \quad (\theta_2, v) = \delta(svc, v_1, \dots, v_n, \theta_1)}{\theta_1 : \text{\$builtin } svc(v_1, \dots, v_n) \rightarrow^\theta \theta_2 : v} \text{ (SVC-UPDATE)}$$

Fig. 15 Interface with the δ function

While there may be some services or operators for which it could be possible to extract a simple operational semantics, for purposes of cohesion we apply the same approach to the semantics of all of them.

4.3.1 Calling a Service

There are two different types of calls to services: those that are internal of other services and operators (e.g., the iterator function `next` calls the `type` service), and those that are invoked by the user. This distinction is important because it is possible to change the binding of a service but only in the latter situation:

Example 12 (Overriding a built-in service.)

```
print(type({})) -->> table
type = function() return 'foo' end
print(type({})) -->> foo
next(1) -->> error: bad argument (table expected, got number)
```

As we can see from this example, the `next` function was unaffected by the change to the `type` function, while the user is now unable to call Lua’s built-in service.

From the point of view of the formalization, we introduce a new expression **\$builtin** which simply acts as a dispatch to the δ function (Fig. 15). Note that we cannot simply let a program refer directly to the δ function: δ is a meta-function, and it does not belong to the programming language. Therefore, we need an interface *within the language*. Then, prior to the execution of a program, the `_ENV` table is populated with functions from the standard library as wrappers to the **\$builtin** constructor.

The semantics of the new expression is simple: it just reduces to an application of δ , providing the necessary arguments and processing its output when needed. We distinguish three cases, according to the semantics of the service: if it does not use the store (SVC-NOSTATE); if it *inspects* the objects store (SVC-CONST); or if it *updates* the object store

$$\begin{aligned}
& \text{indexmeta}(\text{tid}, v, \theta) \doteq \delta(\text{rawget}, \text{tid}, v, \theta) \\
& \text{protmeta}(\text{tid}, \theta) \doteq \text{indexmeta}(\text{tid}, \text{"_metatable"}, \theta) \\
& \text{prot?}(\text{tid}, \theta) \doteq \text{protmeta}(\text{tid}, \theta) \neq \text{nil} \\
\\
& \delta(\text{setmetatable}, \text{tid}, v, \theta_1) = (\text{tid}, \theta_2) \quad \begin{array}{l} \text{if } \delta(\text{type}, v) \in \{\text{"table"}, \text{"nil"}\} \\ \wedge \neg \text{prot?}(\text{tid}, \theta_1) \\ \wedge \theta_2 = \theta_1[\text{tid} := (\theta_1(\text{tid})(1), v)] \end{array} \\
\\
& \delta(\text{setmetatable}, v_1, v_2, \theta) = (\delta(\text{error}, \text{"..."}), \theta) \quad \begin{array}{l} \text{if } \delta(\text{type}, v_2) \notin \{\text{"table"}, \text{"nil"}\} \\ \vee \text{prot?}(\text{tid}, \theta_1) \\ \vee \delta(\text{type}, v_1) \neq \text{"table"} \end{array} \\
\\
& \delta(\text{getmetatable}, \text{tid}, \theta) = \begin{cases} \text{protmeta}(\text{tid}, \theta) & \text{if } \text{prot?}(\text{tid}, \theta) \\ \theta(\text{tid})(2) & \text{otherwise} \end{cases} \\
\\
& \delta(\text{getmetatable}, v, \theta) = \begin{cases} \text{nil} & \text{if } \text{tid}_{\delta(\text{type}, v)} \notin \text{dom}(\theta) \\ \text{tid}_{\delta(\text{type}, v)} & \text{if } \text{tid}_{\delta(\text{type}, v)} \in \text{dom}(\theta) \\ & \wedge \neg \text{prot?}(\text{tid}_{\delta(\text{type}, v)}, \theta) \\ \text{protmeta}(\text{tid}_{\delta(\text{type}, v)}, \theta) & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 16 Primitives for setting and getting metatables

(SVC- UPDATE).⁸ Each rule lists some of the services that fit into each category. The actual list of services includes almost all the built-in basic functions of the Lua language, together with several services from the libraries `math`, `string`, and `table`.

4.4 Metatables

We have seen in repeated occasions how several operations required the intervention of metatables, usually in special situations, like trying to add two values that are not numbers nor strings. Metatables are plain tables that let programmers specify *metamethods* for resolving such situations, associating specific string keys with—typically—functions. For instance, `__add` is the key for associating a handler to the addition operator.

We first explain how to obtain and set metatables. Then, we show with an example how the metatable mechanism triggers once a labeled term is found.

4.4.1 Setting and Getting Metatables

For each Lua type except tables (and userdata, not included in our model), it is only possible to define one single metatable. That is, for example, the type of numbers has one metatable ruling every possible interaction with numbers. For tables, instead, it is possible to define a metatable per value. To this effect, the basic library provides two primitives, `setmetatable`, to set the metatable of a given table, and `getmetatable`, to obtain the metatable associated with a given value (of any type). Being services of the basic library, we provide a declarative description of their behavior using δ . Their semantics is shown in Fig. 16.

⁸ We do not include the value store since our subset implemented so far does not require it, but this may change in the future.

The first equation of `setmetatable` shows the normal behavior: we are setting v as the metatable of tid , and v is itself a table or **nil**; and tid is not *protected*. If these conditions apply, the output of `setmetatable` is the identifier of the table together with a new object store, resulting from updating θ_1 to include v as the second component of the value stored for tid . If v is **nil**, this has the effect of removing any previous association between the table and a metatable.

If a metatable has a field having key "`__metatable`", it is said to be protected. The concept is specified by the predicate `prot?`. Semantically, a protected metatable cannot be overridden.

The second equation of `setmetatable` deals with different erroneous situations: trying to set a metatable for a value which is not a table; trying to set, as a metatable, a value which is not a table or **nil**; or trying to modify a protected metatable. In any of these cases, an error is issued. We write ... to abbreviate the message, which depends on the condition that produced the failure.

The first equation for `getmetatable` shows its normal behavior: when applied over a table value tid , it either returns the value set in its "`__metatable`" field, if present, or the second component in the θ store associated with tid .

More interestingly is the situation when we want to obtain a metatable for a non-table value (the second equation of `getmetatable`). As it turns out, for values such as numbers or strings, it is possible to set the metatable by using the C API or, in the official interpreter, by calling the **debug.setmetatable** function. In our model, we assume that, if set, the metatables for these types are identified by fixed tables' ids: for a given value v , we assume that its metatable must be identified with $tid_{\delta(\text{type}, v)}$. Then, determining if such metatable is set reduces to ask if $tid_{\delta(\text{type}, v)} \in \theta$. Note that such metatables might be protected too.

4.4.2 Resolving Labeled Terms (Excerpt)

In this section, we introduce a new relation $\rightarrow^{\text{meta}}$ that takes a labeled term and acts according to the case that originated the exceptional situation. We do not show all the rules that apply to each different label; instead, we focus on one example rich enough to show how this mechanism works.

Figure 17 describes the $\rightarrow^{\text{meta}}$ relation for field updates. The inner workings of the metatables mechanism is handled by the meta-function `new_index`, while $\rightarrow^{\text{meta}}$ serves as an interface with it (M-UPD). Recall from Sect. 4.2.1 that we have two situations, each marked by its distinctive label: either the update is attempted on a value that is not a table (NONTABLE), or the value is a table, yet the key does not belong to it (WRONGKEY). In either case, if there is a function associated with the key "`__newindex`" in the metatable of the object (first equation of `new_index`), the function is called with the arguments of the update. If, instead, the handler is not a function, nor **nil**, the update is performed on the handler (second equation).

If the handler is **nil**, we need to distinguish according to the label. For WRONGKEY, we try creating the field with **rawset**. If it succeeds (third equation), we return the modified θ store and the skip statement. If it fails (fourth equation), because the key is **nil** or **nan**, we propagate the error. Note that we need to inspect the returned error object **\$err**—the result of $\delta(\text{error})$ —which is introduced later in Sect. 4.5. Finally, if the label is NONTABLE, we just produce an error (last equation of `new_index`).

$$\begin{aligned}
 \text{new_index}(\theta : \langle v_1 [v_2] = v_3 \rangle_{\text{LABEL}}) &= \theta : \mathbf{\$statFcall} \ v_4 \ (v_1, v_2, v_3) \\
 &\quad \text{if} \ \begin{cases} \text{LABEL} \in \{\text{WRONGKEY}, \text{NONTABLE}\} \\ v_4 = \text{indexmeta}(v_1, \text{"_newindex"}, \theta) \\ \delta(\text{type}, v_4) = \text{"function"} \end{cases} \\
 \\
 \text{new_index}(\theta : \langle v_1 [v_2] = v_3 \rangle_{\text{LABEL}}) &= \theta : v_4 [v_2] = v_3 \\
 &\quad \text{if} \ \begin{cases} \text{LABEL} \in \{\text{WRONGKEY}, \text{NONTABLE}\} \\ v_4 = \text{indexmeta}(v_1, \text{"_newindex"}, \theta) \\ v_4 \neq \mathbf{nil} \\ \delta(\text{type}, v_4) \neq \text{"function"} \end{cases} \\
 \\
 \text{new_index}(\theta_1 : \langle \text{tid} [v_1] = v_2 \rangle_{\text{WRONGKEY}}) &= \theta_2 : ; \\
 &\quad \text{if} \ \begin{cases} \text{indexmeta}(\text{tid}, \text{"_newindex"}, \theta_1) = \mathbf{nil} \\ (\theta_2, \text{tid}) = \delta(\text{rawset}, \text{tid}, v_1, v_2, \theta_1) \end{cases} \\
 \\
 \text{new_index}(\theta_1 : \langle \text{tid} [v_1] = v_2 \rangle_{\text{WRONGKEY}}) &= \theta_2 : \mathbf{\$err...} \\
 &\quad \text{if} \ \begin{cases} \text{indexmeta}(\text{tid}, \text{"_newindex"}, \theta_1) = \mathbf{nil} \\ (\theta_2, \mathbf{\$err...}) = \delta(\text{rawset}, \text{tid}, v_1, v_2, \theta_1) \end{cases} \\
 \\
 \text{new_index}(\theta_1 : \langle \text{tid} [v_1] = v_2 \rangle_{\text{NONTABLE}}) &= \theta_2 : \delta(\text{error}, \text{"error updating..."}) \\
 &\quad \text{if} \ \text{indexmeta}(\text{tid}, \text{"_newindex"}, \theta_1) = \mathbf{nil} \\
 \\
 \frac{\text{new_index}(\theta_1 : \langle v_1 [v_2] = v_3 \rangle_{\text{LABEL}} = \theta_2 : s}{\theta_1 : \langle v_1 [v_2] = v_3 \rangle_{\text{LABEL}} \xrightarrow{\text{meta}} \theta_2 : s} &\quad (\text{M-UPD})
 \end{aligned}$$

Fig. 17 Metatable mechanism for field update

4.5 Semantics of Programs and Error Handling

We concluded the presentation of the different relations used to provide semantics to Lua. Following Sect. 3, we must now provide meaning to entire programs, defining a new relation \mapsto that includes execution order and forwards to the specific relation the next piece of the program to be executed. Since its definition is mostly an extension of the relation shown in Fig. 3, but now adding also the θ context, we will omit these definitions for brevity. More interestingly, we develop in this section the treatment of errors in Lua, whose propagation requires its formulation at the level of the \mapsto relation, for soundness reasons.

4.5.1 Errors and Protected Mode

Lua provides mechanisms to generate and catch errors. Errors are represented with *error objects* or *error messages*. They can be explicitly generated by user code, using the function **error**, and it is possible to attach information about the error using any Lua value. The semantics of error objects consist, essentially, in aborting the execution of the program where they are generated.

In order to catch an error, the user is provided with two functions, **pcall** and **xpcall**, the former being a specialization of the latter. They allow to execute a function under a *protected mode*: any error thrown by the function is caught, avoiding error propagation. For

an erroneous function call, **pcall** returns **false** and the error object. Otherwise, it returns **true** together with the values returned by the function call. **xpcall** behaves as **pcall**, but instead of returning the error object in case of error, it forwards it to a user-provided handler function to continue the execution.

4.5.2 Formalization

In formalizing error propagation, we need to take into account that it is a context-sensitive computation, requiring to look at the whole remaining program in order to decide about program termination. Therefore, we need to be cautious if we want to obtain a deterministic model. To understand this, let us consider a program of the form $E \llbracket \mathbf{\$err} \ v \rrbracket$, where E does not contain a protected mode. What we want for such a program is to completely halt in one step, ending as $\mathbf{\$err} \ v$. This suggests to include a step:

$$E' \llbracket \mathbf{\$err} \ v \rrbracket \rightarrow^{s/e} \mathbf{\$err} \ v$$

But, since \mapsto closes over $\rightarrow^{s/e}$ as follows:

$$\frac{t \rightarrow^{s/e} t'}{\sigma : \theta : E'' \llbracket t \rrbracket \mapsto \sigma : \theta : E'' \llbracket t' \rrbracket}$$

we have several possible ways of propagating an error, depending on the decomposition of context E between contexts E' and E'' , offending the intended determinism of the semantics.

The common workaround to this problem is to propagate errors directly from \mapsto . In absence of a protected mode to catch an error, it will make the whole program to halt.

Figure 18 shows the new constructions added to the model for formalizing errors and protected mode. Error objects will be denoted with the construction $\mathbf{\$err} \ v$, which allows for the inclusion of any Lua value as information attached to the error. $(e)_{\text{PROTMD}}^v$ and $(E)_{\text{PROTMD}}^v$ represent an expression or a context that must run in protected mode, where v is the error handler passed to **xpcall**. Treating function calls in protected mode just as expressions will simplify the semantics, as shown below. We also include the same protected mode labels but without the handlers, which has a particular meaning to be shown below. Error propagation will be decided in terms of the presence or absence of a protected mode. To define the absence, we use a new context E_{np} that is a subset of E but without protected modes (omitted for brevity).

With the given constructions, we can formalize what it means for an error to propagate or to be caught in protected mode. The first rule (E- TERMINATION) aborts the whole program if there is no protected mode around the error object. This is the only rule from error handling that is defined in the \mapsto relation.

The remaining rules are defined in the $\rightarrow^{s/e}$ relation, and concerns the execution of code inside a protected mode. If the code executed in protected mode runs until completion without an error (E- PROTTRUE), then its values are appended to the tuple starting with **true**. Instead, if there is an error, we have two situations: if the handler is a function, then it is executed within a protected mode (without a new handler, E- PROTHANDLER). Note how the call is performed within parenthesis, to take just the first value. Also, note how it is not possible to call the handler if it is not a function (E- PROTHANDLERERR), that is, this scenario cannot be meta-programmed through metatables. In this case, we return a tuple starting with **false** and a descriptive error message.

Once the code of a handler is executed in protected mode (denoted by a PROTMD context without a handler), we find two options: in a successful execution of the handler, its

Grammar

$s ::= \dots \mid \mathbf{\$err} v$
 $e ::= \dots \mid \mathbf{\$err} v \mid \langle e \rangle_{\text{PROTMD}}^v \mid \langle e \rangle_{\text{PROTMD}}$
 $E ::= \dots \mid \langle E \rangle_{\text{PROTMD}}^v \mid \langle E \rangle_{\text{PROTMD}}$

Semantics

$$\frac{E_{np} \neq []}{\sigma : \theta : E_{np} [\mathbf{\$err} v] \mapsto \sigma : \theta : \mathbf{\$err} v} \text{ (E-TERMINATION)}$$

$$\langle \langle v, \dots \rangle \rangle_{\text{PROTMD}}^v \xrightarrow{s/e} \langle \mathbf{true}, v, \dots \rangle \text{ (E-PROTTRUE)}$$

$$\frac{\delta(\text{type}, v_2) = \text{"function"}}{\langle E_{np} [\mathbf{\$err} v_1] \rangle_{\text{PROTMD}}^{v_2} \xrightarrow{s/e} \langle (v_2(v_1)) \rangle_{\text{PROTMD}}} \text{ (E-PROTHANDLER)}$$

$$\frac{\delta(\text{type}, v_2) \neq \text{"function"}}{\langle E_{np} [\mathbf{\$err} v_1] \rangle_{\text{PROTMD}}^{v_2} \xrightarrow{s/e} \langle \mathbf{false}, \text{"error"} \rangle} \text{ (E-PROTHANDLERERR)}$$

$$\langle v \rangle_{\text{PROTMD}} \xrightarrow{s/e} \langle \mathbf{false}, v \rangle \text{ (E-PROTFALSE)}$$

$$\langle E_{np} [\mathbf{\$err} v_1] \rangle_{\text{PROTMD}} \xrightarrow{s/e} \langle \mathbf{false}, \text{"error"} \rangle \text{ (PROTERR)}$$

$$\delta(\mathbf{xpcall}, v_1, v_2, v_3, \dots) = \langle v_2 (v_3, \dots) \rangle_{\text{PROTMD}}^{v_2}$$

Fig. 18 Error object and protected mode

only returned value is appended to a tuple starting with **false** (to signal that a handled error occurred, E-PROTFALSE), and, in an erroneous execution of the handler, an error is returned (E-PROTERR), exactly as in E-PROTHANDLERERR. Finally, note that, by considering function calls in protected mode just as expression, we simplify the last step of the mechanism: we do not need to reason if a given tuple must be discarded or not. We are into an expression, so the tuples are always returned.

Here concludes our tour of the semantics of the language, proceeding to state its properties and discuss its mechanization in the coming sections.

5 Properties of the Semantics

In this section, we state two fundamental properties of our model: *soundness* (as the traditional combination of *progress* and *preservation*) and determinism. Their statements were tested using a lightweight approach enabled by Redex [17], in which we specify a property and use *random testing* to try to falsify it. Together with the successful execution of the code from the test suite of Lua, they help build trust in the model. The details of the experience with

Contexts for well-formedness of terms (excerpt)

$C ::= [] \mid \mathbf{while} \ e \ \mathbf{then} \ C \mid (\ C \)_{\mathbf{BREAK}} \mid \mathbf{local} \ x, \dots = e, \dots \ \mathbf{in} \ C$
 $\mid \mathbf{function} \ \ell(x, \dots) \ C \ \mathbf{end} \mid \dots$

Well-formedness of terms (selected rules)

$$\frac{C \vdash_{wft} \sigma : \theta : e \quad C \vdash_{wft} \sigma : \theta : s_1 \quad C \vdash_{wft} \sigma : \theta : s_2}{C \vdash_{wft} \sigma : \theta : \mathbf{if} \ e \ \mathbf{then} \ s_1 \ \mathbf{else} \ s_2} \text{ (WF-IF)}$$

$$\frac{r \in \text{dom}(\sigma)}{C \vdash_{wft} \sigma : \theta : r} \text{ (WF-REF)}$$

$$C' \llbracket \mathbf{local} \ x_1, \dots, x, x_2, \dots = e, \dots \ \mathbf{in} \ C'' \rrbracket \vdash_{wft} \sigma : \theta : x \text{ (WF-VARLOCAL)}$$

$$C' \llbracket \mathbf{function} \ x_1 \ (x_2, \dots, x, x_3, \dots) \ C'' \rrbracket \vdash_{wft} \sigma : \theta : x \text{ (WF-VARFUN)}$$

$$\frac{C' \neq C'' \llbracket \mathbf{function} \ \ell \ (x, \dots) \ C'' \ \mathbf{end} \rrbracket}{C \llbracket \mathbf{while} \ e \ \mathbf{do} \ C' \ \mathbf{end} \rrbracket \vdash_{wft} \sigma : \theta : \mathbf{break}} \text{ (WF-BREAKWHILE)}$$

$$\frac{C' \neq C'' \llbracket \mathbf{function} \ \ell \ (x, \dots) \ C'' \ \mathbf{end} \rrbracket}{C \llbracket (\ C' \)_{\mathbf{BREAK}} \rrbracket \vdash_{wft} \sigma : \theta : \mathbf{break}} \text{ (WF-BREAKUNFOLDED)}$$

$$\frac{C \vdash_{wft} \sigma : \theta : e \dots \quad C \llbracket \mathbf{local} \ x, \dots = e, \dots \ \mathbf{in} \ [] \rrbracket \vdash_{wft} \sigma : \theta : s}{C \vdash_{wft} \sigma : \theta : \mathbf{local} \ x, \dots = e, \dots \ \mathbf{in} \ s} \text{ (WF-LOCAL)}$$

$$\frac{C \vdash_{wft} \sigma : \theta : e \quad C \llbracket \mathbf{while} \ e \ \mathbf{do} \ [] \rrbracket \vdash_{wft} \sigma : \theta : s}{C \vdash_{wft} \sigma : \theta : \mathbf{while} \ e \ \mathbf{do} \ s} \text{ (WF-WHILE)}$$

$$\frac{C \vdash_{wft} \sigma : \theta : v_2 \quad C \vdash_{wft} \sigma : \theta : v_1 \quad C \vdash_{wft} \sigma : \theta : v_3 \quad v_1 \in \text{dom}(\theta) \quad v_2 \notin \text{dom}(\theta(v_1)(1))}{C \vdash_{wft} \sigma : \theta : (\ v_1[v_2] = v_3 \)_{\mathbf{WRONGKEY}}} \text{ (WF-WRONGKEY)}$$

Well-formedness of configurations

$$C \vdash_{wfc} \sigma : \theta : t \doteq \forall r \in \text{dom}(\sigma), C \vdash_{wft} \sigma : \theta : \sigma(r) \tag{1}$$

$$\wedge \forall \text{tid} \in \text{dom}(\theta), C \vdash_{wft} \sigma : \theta : \theta(\text{tid})(1) \tag{2}$$

$$\wedge C \vdash_{wft} \sigma : \theta : \theta(\text{tid})(2) \tag{3}$$

$$\wedge C \vdash_{wft} \sigma : \theta : t \tag{4}$$

Fig. 19 Well-formedness of configurations (excerpt)

the mechanization are presented in Sect. 6, but it is worth mentioning that we were able to catch, with little effort, several ambiguities and ill-defined cases.

When we say *preservation*, we say that a *well-formed configuration* that can perform a step results in another well-formed configuration. As for *progress*, we say that a well-formed configuration represents the *final computation* of a program, or it can perform a step of computation. In essence, together they allow us to exclude *stuck terms*, computations that are incorrect from the point of view of the semantics, like a program with references not allocated in the store. We explicit all these terms next.

The most important ingredient is the predicate that captures well-formedness of configurations (bottom of Fig. 19), which states the conditions that a given configuration $\sigma : \theta : t$ must hold to be well formed, for the given stores σ and θ , and term t (a statement or expression). This predicate is parametrized by a new context C , whose excerpt can be found at the top of Fig. 19, and which captures the contextual information required to check the well-formedness of t . In essence, C keeps track of local variables, while loops, labeled terms, and function definitions. Using just term contexts to keep track of contextual information allows us to include every information needed to reason about well-formedness in a single object that is also just a kind of term context. That is, we do not need to define new constructs with the corresponding operations to manipulate them. We just reuse the language constructions, and the operations for matching and manipulation of contexts. For example, we can maintain not only information about variables in scope, but also occurrences of while loops or labeled terms, with the same kind of construction, and ask about that information by means of plain pattern matching.

A configuration $\sigma : \theta : t$ is well formed for a context C , written $C \vdash_{wf_C} \sigma : \theta : t$, if every term of σ and every table in θ is well formed, and also t is. When the context is $\llbracket \ \rrbracket$, we omit it and simply write $\vdash_{wf_C} \sigma : \theta : t$. In order to check when a given term is well formed (w.r.t. the well-formed stores), we use a new relation \vdash_{wf_t} and show only a representative set of rules.

\vdash_{wf_t} is defined recursively over the structure of the term t , as can be seen, for example, in WF- IF. Some rules require checking side conditions, like WF- REF, in which the reference must exist in the corresponding store, in order to check for the absence of *free occurrences* of references that would render the corresponding term *stuck*. Additionally, some rules require specific contextual information: for a variable x , we need to check that the context contains a **local** or **function** signature binding x (WF- VARLOCAL and WF- VARFUN, respectively). Similarly, a **break** can only occur if the context contains a yet-to-be-executed **while** (WF- BREAKWHILE) or an already-executed one (WF- BREAKUNFOLDED). In both cases, we also need to check that the **break** statement is actually *immediately* inside the loop or labeled term: this is that the statement is not inside another construction surrounding it, like a function definition. The contextual information is extended each time we check certain constructs. For instance, WF- LOCAL shows how local variables are inserted in the context C prior to checking the inner statement s . Similarly, WF- WHILE shows how the **while** construct is inserted in the context C in order to check the body of the loop. Finally, for labeled expressions or statements representing situations that must be handled by metatables, we need to make sure that the conditions that created the labeling holds. For instance, for WRONGKEY, we need to check that the key indeed does not exist (WF- WRONGKEY).

Now that we have a precise meaning of what constitutes a well-formed configuration, and prior to state the main property of the formalization, we enunciate what *final computations* are

Definition 1 We say that a statement s is *final* if it is one of the following, for some value v :

- (a) **\$err** v
- (b) **return** v, \dots
- (c) **;**

Lemma 1 (Soundness of \vdash_{wf_C}) For given stores σ, θ , statement s , if $\vdash_{wf_C} \sigma : \theta : s$, then one of the following situations holds:

- (a) s is final.
- (b) There exists a configuration $\sigma' : \theta' : s'$ such that $\sigma : \theta : s \mapsto \sigma' : \theta' : s'$ and $\vdash_{wf_C} \sigma' : \theta' : s'$.

Note that in order to prove this and the following statements, we require a clear formalization of the operations of matching and decomposition. Then, we could prove the soundness of \vdash_{wfc} by induction on the structure of the proof of $\vdash_{wfc} \sigma : \theta : s$.

The previous property does not state that the step is unique. In order to have deterministic semantics there should always be at most one redex pointed by an evaluation context, into a given term:

Lemma 2 (Unique decomposition) *Given stores σ , θ , and statement s , such that $\vdash_{wfc} \sigma : \theta : s$, then one of the following situations holds:*

- (a) s is final.
- (b) *There exist only one evaluation context E and term t such that s matches against the pattern $E \llbracket t \rrbracket$, in such a way that only one of the following is a redex:*
 - (1) t
 - (2) $\sigma : t$
 - (3) $\theta : t$
 - (4) $\sigma : \theta : t$

In order to prove it, given that we are looking for the next redex into a given term, we could rely on the information that we have in the proof of $\vdash_{wfc} \sigma : \theta : s$. It checks the side conditions of every semantics rule, in order to guarantee that a given term is not stuck: that is, if it is not a result, then it is a redex. Then, we could prove this statement by means of induction on the structure of the proof of $\vdash_{wfc} \sigma : \theta : s$.

Additionally, the proof will depend on a correct correspondence between evaluation contexts and semantics relations. That is, that each evaluation context should point to just one term that is the next redex, according to the semantics relations.

From these properties, we infer two relevant corollaries. The first one is a direct consequence of the previous lemmas: there are no stuck terms, and the semantics is deterministic. Hence:

Corollary 1 *A well-formed configuration runs until completion, or loops forever. In other words, every erroneous situation is handled properly by the error handling mechanism:*

For every configuration $\sigma : \theta : s$, such that $\vdash_{wfc} \sigma : \theta : s$, just one of the following situations holds:

- (a) *The execution diverges.*
- (b) *The execution ends with a configuration $\sigma' : \theta' : s'$, such that $\vdash_{wfc} \sigma' : \theta' : s'$ and s' is final.*

Finally, the following corollary of soundness of \vdash_{wfc} justifies the representation of closures explained in Sect. 4.2.3:

Corollary 2 (No free variables introduced) *For a given closed term s (i.e., without unbounded occurrences of variables' identifiers), and well-formed configuration $\sigma : \theta : s$, if there exists σ' and θ' such that $\sigma : \theta : s \mapsto \sigma' : \theta' : s'$, then s' is a closed term.*

Note that, otherwise, the resulting configuration $\sigma' : \theta' : s'$ could not be considered well formed, as the soundness of \vdash_{wfc} guarantees. This means that, once the focus of execution has reached to a closure, every external local variable in its body has been replaced by the corresponding reference, i.e., the environment of the closure is completely embedded into its body, avoiding stuck terms.

6 Mechanization

In this section, we discuss the mechanization in Redex of the model presented in this paper. We begin by explaining the differences between the formal model and its mechanization and then discuss the results of executing the tests from the test suite and of randomly testing fundamental properties.

6.1 Correspondence with the Formal Model

During the process of formalizing the language, we went back-and-forth between our pen-and-paper formalization and its mechanization in Redex. This tool was of great help, as the resulting mechanized formalization can be executed and traced; therefore, we could recognize problems in our first attempts at formalizing the language, experiment with new ideas before landing them on the formalization and testing the mechanization against the test suite for the official interpreter, prior to verifying the desired properties.

Given that we produced two documents—the formalization on paper and its mechanization—we did not follow closely the original philosophy of Redex, consisting in considering *semantic models as software artifacts* [17]. Although it is indeed an interesting approach, with a real potential to optimize the time spent in a formalization effort, we still need to observe that reading the mechanization alone is at times a difficult task: some concepts from our model are not possible to express directly in Redex, forcing us to obtain a mechanization that loses some abstraction and clarity. This occurs, for example, when dealing with concepts specified in terms of a quantification over an arbitrary predicate: Lua function calls, operations that filter fields of a given table, etc.

This said, the formalization presented in this paper and its mechanization are remarkably similar. There are, however, four groups of differences that are worth mentioning.

6.1.1 Particularities of the Tool

Redex fully supports the set of concepts of reduction semantics with evaluation contexts; namely, the definition of languages, evaluation contexts, meta-functions, and reduction rules. In particular, reduction rules are described as relations over *terms* of a defined language. In practice, this means that we must represent some semantics elements of our model as phrases of a language. For instance, we presented the stores as partial functions in Sect. 3, but they must be defined as terms in the mechanization. This is observable for instance in the rule for updating a variable:

$$\frac{\sigma' = \sigma[r := v]}{\sigma : r = v \rightarrow^\sigma \sigma' ;} \text{ (LOCAL-ASSGN)}$$

The same rule in Redex is written as follows:

```
[-->σ ((vsp_1 ... (r v_1) vsp_2 ...) : (r = v_2))
      ((vsp_1 ... (r v_2) vsp_2 ...) : \;)
  Local-Assgn]
```

where the σ store is now described in terms of a phrase, and `vsp` is a new syntactic class for describing a pair containing a reference and a value, something that we did not considered

part of the language in the pen-and-paper formalization.⁹ Also, the substitution of the value associated with the reference r by the value v_2 , is now performed using pattern matching and a *term template* [16], which builds a new σ storage that differs from the first one just in the value associated with the reference r : now, v_1 is replaced by v_2 (a pair of a reference and a value is noted with (rv)).

While in this simple example, the distance is short, in some complex rules, the difference becomes more noticeable.

6.1.2 Particularities of the Use of the Semantics

When using our model to build a tool to check misuses of weak tables [38] (briefly introduced in Sect. 7), we had to introduce two changes in order to better emulate the behavior of a garbage collector. First, we change the introduction of local variables (*e.g.*, LOCAL-DECL in Sect. 4.2.2) and the instantiation of formal parameters of functions (*e.g.*, E-CALL in Sect. 4.2.3). We do it in a way that the variables' references that are still in scope are maintained (in some structure) through the whole execution of the code in their scope.

In a common implementation of a compiler of a programming language with reachability-based garbage collection, a variable that is in scope is always considered reachable, regardless of its actual reachability from what is left to be computed. This is so because of the way in which the root set of references is calculated (for example, just by looking at the stack of a function call).¹⁰ On the other hand, in the model presented so far, once that a variable is not reachable from what is left to be computed in its scope, we would consider it unreachable. It does not matter if the code that belongs to its scope is still under execution. This difference with real implementations naturally introduces some observable effects in the context of a programming language with interfaces to its garbage collector, as it happens with Lua 5.2.

In order to better approximate the behavior of a real implementation, we keep the references that correspond to introduced variables within a **local** statement or function call, annexed to the code in which they are replaced. This allows for a simple representation of the set of references of variables that are in scope, regardless of their presence in what remains to be computed. Second, and for similar reasons, closures are manipulated through references, analogous to the manipulation of tables. In fact, this is closer as to how the official implementation of Lua works, and allows us to have a model that better represents the common phenomena associated with garbage collection.

6.1.3 Particularities of the Official Interpreter

In order to stay close to the official interpreter and its corresponding test suite, we had to enrich our specification with implementation details that might not be strictly considered as belonging to the semantics of the language. In many cases, these were just corner cases of the services of the language and its standard library (discussed in Sect. 6.3). But a more interesting example is how to deal with mutually recursive references between metatables and handlers, as the following program illustrates

⁹ The language of patterns from Redex is expressive enough to be able to impose, for example, the expected representation invariant for a term that describes a finite functional mapping.

¹⁰ The first set of references from which reachability is computed.

Example 13 (Mutually recursive references in metatables.)

```

1 local a = {}
2 a.__newindex = a
3 setmetatable(a, a)
4 a[1] = 2

```

From the point of view of the semantics of metatables (Sect. 4.4), the previous program is an infinite loop: the instruction in line 4 triggers the metatable mechanism, looking for the metatable of `a` to index it with the key “__newindex” (rule M-UpdNotFun in Fig. 17), obtaining the handler. But this is a itself, as specified in line 2, repeating the assignment in an endless loop.

While the reference manual does not specify a particular way to handle these situations, the official interpreter implements some intelligence that is capable of recognizing these loops. In particular, the execution of the previous program results in an error with the message “loop in settable”.

Given that the previous behavior is tested in the official test suites, we include it into our semantics by maintaining, in the instruction term, the tids to the metatables is used while trying to solve a labeled term. In this way, we can recognize when the same metatable is being indexed repeatedly and raise the appropriate error.

6.1.4 Particularities of the Implementation: The load Service

Like many scripting languages, Lua includes a service to execute an arbitrary program coded in a string. This service, called **load**, has little interest from the point of view of its semantics: given a string `s`, it returns a function that, when called, executes the code written in `s`. The service returns **nil** if there are syntax errors or unbound identifiers in the resulting program.

From the point of view of its mechanization, however, it is interesting to mention that its inclusion was rather trivial: our parser and the desugaring process are both written in Racket, the underlying language of Redex, and Redex itself.

As in [34], we also split desugaring into 2 pieces: a compiler from Lua source code to our language that produces (possible) open terms; and a definition of an evaluation context into which the previously compiled term is plugged. This context contains the standard implementation of the execution environment: a table, containing the set of bindings to library services, which, in our model, are just invocations to the services through the **\$builtin** construction. This is our definition of the original execution environment `_ENV` (see Sect. 2.2). Under normal circumstances, `load` will assume that the global variables in the interpreted code refer to said environment: they are just fields of the table bound to `_ENV`. However, the user can pass, as an argument to `load`, a table that will be used as a replacement to `_ENV`. In that case, global variables will be interpreted as fields of the given table.

6.2 Tests Coverage

From the 25 files present in Lua’s test suite [21], we successfully run 12 files (including garbage collection; see [38]), with varying degrees of coverage (Fig. 20). Each file from the test suite is a sequence of assertions about the expected outcome of the code. Those files and lines not tested are due to the following reasons:

File	Features tested	LOCs (total/tested)	Coverage
calls.lua	functions and calls	221/162	73.3%
closure.lua	closures	198/139*	70.2%
constructs.lua	syntax and short-circuit opts.	237/193*	81.4%
errors.lua	errors	317/137*	43.2%
events.lua	metatables	302/300	99.3%
gc.lua	garbage collection	445/181*	40.1%
locals.lua	local variables and environments	114/75	65.8%
math.lua	numbers and math lib	219/191	87.2%
nextvar.lua	tables, next, and for	355/226*	63.7%
sort.lua	(parts of) table library	133/46	34.9%
strings.lua	strings and string library	233/88*	37.8%
vararg.lua	vararg functions	95/95	100%
Total		2869/1833	63.9%

* These files were slightly modified in order to increase performance and coverage of tests. Counted LOCs do not include extra lines added for these purposes.

Fig. 20 Lua 5.2's test suite coverage

- Language features not covered by our formalization: coroutines, the **goto** statement, some standard library functions (mostly related with file handling), and other standard libraries implemented in C (bit32, coroutine, debug, io, etc.).
- Tests covering implementation details of the interpreter and not relevant to the language's semantics: tail call implementation, manipulation of large tables (not tested for performance reasons), generation of bytecode, and the like.¹¹

In pursuit of reasonable running times for each module of tests, we divided each .lua file in smaller modules. Also, when necessary, we reduced the iterations performed by loops and added explicit calls to the garbage collector (16 LOCs, affecting results in a few other lines). Additionally, while looking for more coverage, we changed a handful of lines in modules errors.lua and strings.lua to remove services that are not yet included. In the case of errors.lua, the lines were in functions used throughout the code to test the error messages. By modifying them, we obtained weaker predicates for testing our formalization of the error mechanism, without testing details of the produced error messages.

With the addition of parallel threads of execution, enabled by modern-day processors, the whole selected fraction of Lua's test suite can take, roughly, 26 min to completion without

¹¹ While we do model tail calls, the actual testing of the mechanism in the official test suite consists in performing several recursive calls, beyond the stack limit, to actually test if the mechanism is being used. Given that performing such amount of recursive calls is not feasible within our mechanization (because of the time required), we do not include such tests.

testing garbage collection, and 1:22 h including it.¹² Therefore, it is possible to add the execution of these tests in CI to ensure that new features and fixes included in the language do not break existing ones.

The accompanying material includes the modified files with relevant information to inspect the changes and run the tests.

6.2.1 Improving Performance Thanks to Semantics-Preserving Garbage Collection

As mentioned, in order to reduce the timings of the test, we introduce explicit calls to our mechanization of Lua's garbage collector. To understand the reasons, consider the execution of the following Lua program:

```
for i = 1, 100 do; end
for i = 1, 100 do; end
```

Without the intervention of the garbage collector, each iteration of both loops generates garbage that will persist in the configurations until the end of the execution.¹³ Given the computational model of reduction semantics—in which at each step, the whole configuration is matched against the (possible) several patterns of the left side of each execution rule—increasing the size of the store impacts negatively on the performance of this task. Hence, if we reduce the size of the store during run time, in a way that preserves the semantics of the program, we obtain some performance benefits without compromising soundness.

For instance, the execution of the previous program takes roughly 2 minutes on the same hardware. By adding a single call to the garbage collector, right after the first loop, the execution time is shortened by more than a half. This said, in smaller configurations the performance of the execution of the garbage collector can also impact negatively. For example, if we consider a modification of the previous program where both loops iterate 10 times, the program that performs garbage collection after the first loop is slower (by a couple of tenths of a second) than the one that runs until the end carrying all the garbage generated by both loops. We understand this as the cost of having to execute our garbage collector, which inherits some of the complexities of Lua's—yet not its efficiency.

Together with the modularization of the code into smaller files, adding explicit calls to the garbage collector allowed us to mitigate the known problems with performance that comes with mechanizing an interpreter in Redex [35], and in a semantic-preserving way [38].

6.3 Random Testing of Properties

Redex's features for implementing formal systems and random testing (`redex-check`[15]) allowed us to mechanize the desired check for well-formedness of programs and to stress test its soundness property, together with the expected determinism of the semantics.

Interestingly, the randomly generated terms showed unexpected errors and omissions in our mechanization that were not caught during testing against the test suite. As it turns out, `redex-check` proceeds completely *uninhibited*, unaware of the semantics [17], providing intricate, hard-to-decipher terms that proved to be useful to uncover:

¹² On a system running Arch Linux updated to May, 2022, Racket v8.3, and 12 parallel processes on an Intel Core i7-10870H CPU @ 4GHz × 8, with 16GB of RAM.

¹³ Remember that, at each iteration, a new local variable *i* is created, to contain the corresponding value in the interval [1, 100].

1. Missing border cases and implicit coercions of the parameters of some library services.
2. Erroneously handled cases of the table constructor, mostly related with the generation of numeric keys.
3. Misinterpretations from our part of the reference manual.
4. Errors in our first attempts at defining a suitable notion of *well-formed configuration* for our semantics.

With respect to 1, these omissions can be understood as behavior that is assumed or which does not belong to the essentials of the semantics. Naturally, the later cases cannot be eluded from our semantics, since we aim at an executable model. As an example, the results of the following invocations of library services cannot be inferred from the reference manual:

Example 14 (Special cases in library services)

```

1 -- sub(s,i,j): substring of s from position i and with length j
2 print(string.sub("abc", 1, -3))    -->> "a"
3
4 -- rep(s,n,sep): repeat string s n times, each copy separated by sep
5 print(string.rep("abc", 1.99999, ""))    -->> "abc"

```

In the call to **string.sub** the third argument is negative, a situation that—while mentioned in the reference manual—its semantics is not actually specified. In the call to **string.rep** the parameter *n* is not an integer. In this case, the official implementation takes the floor, although this is not mentioned in the manual.¹⁴

Even though our first attempt at mechanizing the semantics contained several of these errors, it actually managed to pass the tests. We remark that it is not necessarily a problem of the tests of the interpreter, since several services are just wrappers around already tested C's libraries.

In this scenario, using *redex-check*, for testing easily verifiable properties like progress, served as a lightweight approach to stress testing our implementation in search of omissions and errors, with a reasonable time overhead.

6.3.1 Preparation of Test Cases

Naturally, if run *inhibited*, most of the terms generated by *redex-check* are ill-formed, since the grammar does not enforce well-formedness. To alleviate this, it is possible to provide *redex-check* with a preparing function that manipulates the generated term and turns it into a well-formed one, when possible.

Additionally, to obtain a better coverage of the rules, *redex-check* offers the possibility of generating terms following the pattern of the left-hand side of each of the rules. However, while we obtain full coverage, this selection of patterns is too restrictive: the properties that we are interested in random testing, are predicates quantified over a well-formed configuration $\sigma : \theta : s$, *i.e.*, not just terms that matches against the left-hand side of the semantic rules. Therefore, to increase the confidence on the testing, we generate just a fraction of the total cases following the semantics rules, and ask *redex-check* to generate the remaining terms by just following the pattern $\sigma : \theta : s$.

¹⁴ Note about numeric representations: The official interpreter follows the IEEE 754 standard, while Racket has a richer set of numbers and behaviors, even including complex numbers. Therefore, it is required to convert to the required representation of numbers in order to correctly emulate the behavior of Lua.

6.3.2 Properties Tested and Results

For every well-formed term generated, we tested progress and preservation of well-formedness, together with the determinism of our semantics. For a run of 50,000 attempts, we obtain, in an average of ten runs, close to 43,000 well-formed configurations that successfully pass the tests, with complete coverage of the rules.

To better understand the high percentage of well-formed configurations, note that our preparation function takes ill-formed terms and tries to turn them into well-formed ones, also attempting to directly replace conflicting sub-terms by (simpler) well-formed ones.

During random testing, the terms produced by `redex-check` may grow in size and complexity after each failed attempt at generating a term. It is possible to specify a rate at which the terms grow in size, by providing a bound to it. According to the documentation of `redex-check`, said bound “*bounds the height of the generated term (measured as the height of its parse tree)*”. We used the default bound growth rate which grows as the base 5 logarithm of the number of previous attempts at generating terms. To get a general idea of the complexity of the terms generated, on a run with 50,000 attempts, the maximum bound used is 6, which may generate configurations with programs consisting of, roughly, 250 statements (each with its expressions), and stores with 30 elements (without counting those added by the preparation function).¹⁵

7 A Use Case: Modeling Garbage Collection

In this section, we present a brief tour to an existing application of the dynamic semantics presented so far. The content is a distilled version of [38].

7.1 Garbage Collection

As we already mentioned, we used our dynamic semantics to model garbage collection (GC) in Lua. In essence, the task reduces to include a non-deterministic semantics rule that performs a GC cycle. Now, in order to reason about its behavior, we need to study it in the context of a complete dynamic semantics and see how it interacts with the rest of the semantics model. For plain GC, *i.e.*, without considering the interfaces with the garbage collector, we expect each collection step to preserve the semantics of the program. This amounts to *observing* the same behavior of the program before and after the GC step.

A more interesting situation arises when we include the interfaces with the garbage collector: *weak tables* and *finalizers*. Weak tables are tables whose keys and/or values are referred by *weak references*. These references are not taken into account by the garbage collector when determining reachability and can, therefore, be collected at any time. Finalizers are similar, but not equal, to *destructors* in languages with explicitly managed memory. Such interfaces might introduce observably diverging results. Our work enables the development and certification of static analyzers for ruling out potentially dangerous non-determinism from a program.

¹⁵ To obtain these measures we added custom Racket code to our module `Tests/RandomTesting/soundness/soundness_rand_test.rkt`. We do not include it in the provided source code, though the code reduces to counting well-formed configurations and recognizing the biggest configurations generated.

In this section, we first present **LuaSafe** (Sect. 7.2), a tool that uses the modeled semantics of weak tables to ensure deterministic execution of programs, and then we take a sneak peak at how we model (simple) garbage collection (Sect. 7.3).

7.2 LuaSafe in a Nutshell

LuaSafe is a prototype static analyzer that aims at the detection of misuses of weak tables that could lead to non-deterministic behavior. While the general problem is known to be undecidable [3], we propose an approximation to the solution by combining techniques from statics semantics (type inference, type checking, and data-flow analysis) together with weak tables' semantics.

For a given Lua program p , we say that p is *gc-safe* if it exhibits a deterministic behavior under \mapsto extended with GC and its interfaces. Then, we denote with P_{safe} the set of gc-safe programs. In our approach, we aim at taking a user program and trying our best to guess if it belongs to P_{safe} , without asking the user for modifications of the program or to use weak tables according to some particular idioms, as proposed in [3].

In order to achieve this, **LuaSafe** proceeds in a series of steps, including a type system which makes heavy use of the formal tools required to understand Lua's GC. In the coming section we make a dip at the underlying formalism.

7.3 Modeling Garbage Collection

The purpose of GC is to remove from memory (the store) information that will not be used by the remaining computations of the program. One of the simplest and commonly used approaches to find such information is based on the notion of *reachability* (e.g., [6]). The idea is simple: given the set of references that literally occur in the program (the *root set*), it must be the case that any *information* (e.g., value in a store) that may be used by the program must be *reachable* from that set. Conversely, any *binding* (a reference with its value) in the store that cannot be reached from the root set, will not be accessible from the program and, therefore, can be safely removed as it will not be needed in the remaining computations of the program.

Lua implements two reachability-based GC strategies: a *mark-and-sweep* collector (the default) and a *generational collector*. In this section, we will provide a specification for the behavior of a typical reachability-based GC that encompasses the essential details of the behavior of the two algorithms included in Lua and any other based on reachability.

In the context of this work, those values which are not reachable will be called *garbage*. This notion, sufficient to model Lua's GC, is purely syntactic: it will take into account just the literal occurrence of references in the program, or their reachability from this set of references that occur literally, to determine if a given value is garbage or not. In contrast, there are approaches, to identify garbage, where also the semantics of the program may be taken into account (e.g., [3]).

To formally capture the notion of garbage, it will be easier to begin with the definition of reachable references. The only difference worth to mention, in comparison with common definitions found in the literature [6, 19], is the inclusion of metatables: a metatable of a reachable table is considered reachable, so a reachability *path*, that is, a path between a reference and the root set, might also go through a metatable.

Informally, a location (value reference or an identifier) will be reachable with respect to a given term t , and corresponding stores, if one of the following conditions hold:

- The location occurs literally in t .
- The location is reachable from the information associated with a reachable location. This includes
 - The location is reachable from the closure associated with a reachable location.
 - The location is reachable from the table associated with a reachable location.
 - The location is reachable from a metatable of a reachable table identifier.

This is formalized in the following definition:

Definition 2 (*Reachability for Simple GC*) We say that a location $l \in r \cup \text{tid}$ is *reachable* in term t , given stores σ and θ , iff:

$$\begin{aligned} \text{reach}(l, t, \sigma, \theta) = & l \in t \vee \\ & (\exists r \in t, \text{reach}(l, \sigma(r), \sigma \setminus r, \theta)) \vee \\ & \exists \text{tid} \in t, (\text{reach}(l, \pi_1(\theta(\text{tid})), \sigma, \theta \setminus \text{tid}) \vee \\ & \text{reach}(l, \pi_2(\theta(\text{tid})), \sigma, \theta \setminus \text{tid})) \end{aligned}$$

The model proposed in [38] manages closures through references, in a similar fashion as with tables. Here, we avoid that aspect in order to keep it simpler and closer to the model presented in this work, keeping in mind that it is only and approximation of the actual notion of reachability we use in the cited work.

We write $l \in t$ to indicate that l occurs literally in term t and write $\gamma \setminus l$ as the store obtained by removing the binding of l in γ . Informally, this predicate states that either l occurs in t , or there is a reference in t such that l is reachable from it.

To avoid cycles generated from mutually recursive definitions, in the stores that would render undefined the preceding predicate, we remove from the stores the bindings already considered. We assume the predicate is false if a given location occurs in t but does not belong to the domain of any of the stores.

Note that for a table tid , we not only check its content ($\pi_1(\theta(\text{tid}))$) but also its metatable ($\pi_2(\theta(\text{tid}))$). That is, a table’s metatable is considered reachable when the table itself is reachable. Observe that, being metatables ordinary tables, they can contain other tables’ ids or even closures, which in turn may have other locations embedded into them. Naturally, if metatables were not taken into account for reachability, we could run straight into the problem of dangling references any time a metatable is recovered from the metatable. Also, note that during the recursive call $\text{reach}(l, \pi_2(\theta(\text{tid})), \sigma, \theta \setminus \text{tid})$, at first it will determine if l is exactly $\pi_2(\theta(\text{tid}))$ (because it asks for $l \in \pi_2(\theta(\text{tid}))$, for $\pi_2(\theta(\text{tid}))$ being either **nil** or a table identifier) and, if not, it will continue with the inspection of the content of the metatable, by dereferencing its id, given that it is not **nil**. Hence, we do not remove $\pi_2(\theta(\text{tid}))$ from θ in the mentioned recursive call.

7.3.1 Specification of a Garbage Collection Cycle

We keep abstract the specification of a cycle of GC in order to accommodate to any implementation of GC:

Definition 3 (*Simple GC cycle*) $\text{gc}(s, \sigma, \theta) = (\sigma_1, \theta_1)$, where:

- $\sigma = \sigma_1 \uplus \sigma_2$
- $\theta = \theta_1 \uplus \theta_2$
- $\forall l \in \text{dom}(\sigma_2) \cup \text{dom}(\theta_2), \neg \text{reach}(l, s, \sigma, \theta)$

We use $\gamma_1 \uplus \gamma_2$ to denote the union of stores with disjoint domains. This specification states that $\text{gc}(s, \sigma, \theta)$ returns two stores, σ_1 and θ_1 , which are a subsets of the stores provided as arguments, σ and θ . We do not specify how these subsets are determined. We just require that the remaining part of the stores (σ_2 and θ_2) do not contain references that are reachable from the program s . Satisfied this condition, it is safe to run code s in the new stores σ_1 and θ_1 , as no dereferencing of a dangling pointer may occur.

Observe that the previous specification does not impose σ_1 and θ_1 to be *maximal*, meaning they might have non-reachable references with respect to s .

Using the previous specification of GC, we can extend our model of Lua with a non-deterministic step of GC:

$$\frac{(\sigma', \theta') = \text{gc}(s, \sigma, \theta) \sigma' \neq \sigma \vee \theta' \neq \theta}{\sigma : \theta : s \mapsto \sigma' : \theta' : s}$$

We require it to actually perform some changes to the stores to ensure progress. This obviously introduces non-determinism: at any time, as long as there is some garbage left, we can choose to collect the garbage or to continue with the execution of the program. But, for the definition provided so far, this non-determinism should not change the behavior of the program: every execution path will eventually lead to the same result. This property will not longer be true when extending GC with the interfaces with the garbage collector present in Lua 5.2. We further develop the model in the cited bibliography.

8 Related Work

As mentioned throughout the text, the major source of inspiration are the line of work by [8, 34, 35]. The first two of these works formalize JavaScript, and the third one formalizes Python. At a broad view, we share with these works several of the characteristics of the semantic models and their corresponding mechanization in Redex, although there are several differences important to note. First, our model includes particularly interesting features, like Lua's mechanism for meta-programming. While similar to Python's metaclass construction, this feature is not covered in [35]. Second, we purposely do not focus on identifying a core set of features, but instead model the language as a whole (approach enabled by Lua's relatively small size). As mentioned in Sect. 6, we have a very thin desugaring process, staying very close to the surface language. Fourth, and related to the previous one, we tested our mechanization using the language's own test suite directly within Redex, instead of writing another interpreter, as done for performance reasons in [8, 35]. This increases the confidence in our mechanization, at the expenses of spending more testing time. Lastly, we pay special attention on the specification and random testing of the model's properties.

By not following the core language approach, we avoided known complexities in the resulting model [2, 29]: verbose desugared code and a reduced confidence on the compliance of the given semantics with respect to the original language's specification. Additionally, maintaining the proximity with the original language paves the way to a mechanization that also Lua developers could use and verify, as their intuition about the language's semantics is better expressed. This goal is shared with WebAssembly, whose semantics is formally specified in its core documentation [36], and introduced in [9].

Another major step in formal semantics for JavaScript is JSCert [2]: a formalization of ES5 in the Coq proof assistant, together with an interpreter extracted from the formalization (JSRef). It presents a big-step semantics for the specification of ES5. In order to gain confidence about the compliance of their formalization with respect to the specification of ES5,

the authors recognize the importance of the revision of their Coq model by people of different areas, ranging from developers of analysis tools for JavaScript, developers of JavaScript VMs and even ECMA authors. In our project, it is our hope to include Lua developers. This may be difficult to achieve over a Coq model because, as the authors from [2] recognize, using proof assistants requires considerable more learning than using, for example, tools specifically designed for mechanizing language specifications. While we pursue in the future the mechanization of proofs of properties of our model, perhaps using Coq, we want to take advantage of the ease of use of Redex.

Besides the aforementioned [8, 34], [29] introduces a small-step operational semantics for the full language specified in ECMA-262 Standard, 3rd Edition, including proofs of several properties of the model. The authors recognize that, by defining the semantics of each construction as described by the ECMA specification, it gives them the greatest likelihood that the model is correct. While this approach resulted in a model that inherited the size and complexities of the language being defined, the experience is interesting for our investigations on Lua, as we are dealing with a smaller, simpler language.

Specific to Lua's semantics, to the best of our knowledge, there is just one other experience for Lua 5.2: [20] presents an operational semantics for Lua, in the style of Featherweight Java [14], i.e., with a strong focus on recognizing a core language. It considers a subset of the features from the ones presented here and provides a reference implementation in Haskell. Lua 5.1's dynamic semantics has been studied in [31], together with type inference, to obtain an optimized compiler.

9 Final Thoughts and Future Work

This project aims at understanding Lua to the point of being able to construct certified tools, being [38] our first step in this direction. Currently, we are at the point where it is possible to analyze realistic programs, perhaps with coroutines and goto being the most salient missing features.

With respect to the mechanization, Redex proved to be a great tool for the semanticist: it is flexible and it is equipped with a well-suited collection of tools. However, as is expected with a dynamic language as Redex, this flexibility is at times a double-edge sword, as it gives path to several simple bugs to crawl in, bugs that are easily avoided with a type system. This said, we ended up eliminating a large amount of bugs by random testing the semantics.

The interpreter resulting from the mechanization is significantly slow, but it is of no surprise taking into account that the whole mechanization takes about 8k LOCs (without counting comments and tests). Some researchers [8, 35] decided to construct another interpreter by hand, testing this interpreter instead of the mechanization in Redex. In our case, we were still able to execute files with hundreds of lines of code within Redex, by splitting the files and/or adding specific calls to the garbage collector.

An interesting possibility to explore is to port and formally prove the properties in a proof assistant. Then, one could extract a verified interpreter based on abstract machines that should, in principle, run faster than executing the model within Redex, that follows the computation model of reduction semantics. Also, this automated step should help in building trust about the correspondence of the obtained interpreter and the original formal semantics. We are currently interested in that venue of research, and leave for future research extending and updating the model to current versions of Lua.

From the missing aspects of the semantics, the most interesting ones from a semantic point of view are **goto** and coroutines. About the former, it is relevant to mention that it is

not a widely used keyword. A code search¹⁶ shows only 37k occurrences of “goto” in about two millions of files, including a significant fraction of spurious occurrences in comments. As such, it seems the effort required to extend the semantics seems unjustified: program configurations has to be enriched to include information about possible jumps and always keep the already-executed code. This is studied in detail in [18] for a C-like language.

Coroutines, on the other hand, are significantly more used (a similar search returns about 164k hits). Luckily, they are already understood from a formal point of view: [32] presents a reduction semantics specifically tailored for expressing Lua’s coroutines, and it seems it could be easily added on top of the concepts already present in our model.

Acknowledgements We thank Dr. Daniel Fridlender and Dr. Fabio Mascarenhas for their useful insights and support during the development of this research. We also thank wholeheartedly the anonymous reviewers, whose attention to detail helped us deliver a significantly improved paper.

Funding This work was funded by the following projects: Consolidar II 33620180101063CB, UNC SECyT (Argentina) PICT D 2017-3315, ANPCyT (Argentina). The first author received a grant from CONICET, Argentina.

Data availability The claims made in the present work are supported by the provided mechanization. With it, it is possible to reproduce the results shown in the paper.

Code Availability The mechanization is publicly available at <https://github.com/Mallku2/luca-gc-redex-model>.

Declarations

Conflict of interest The authors declare that they have no conflict of interest.

References

1. Adobe: Adobe Lightroom®. <https://www.adobe.io/apis/creativecloud/lightroom.html> (2019). Accessed 04 May 2020
2. Bodin, M., Chargueraud, A., Filaretti, D., Gardner, P., Maffei, S., Naudziuniene, D., Schmitt, A., Smith, G.: A trusted mechanised JavaScript specification. In: POPL ’14 (2014)
3. Donnelly, K., Hallett, J.J., Kfoury, A.: Formal semantics of weak references. In: ISMM ’06 Proceedings of the 5th international symposium on Memory management, pp. 126–137 (2006)
4. Felleisen, M.: The calculi of lambda-v-cs conversion: a syntactic theory of control and state in imperative higher-order programming languages. PhD thesis, Indiana University (1987)
5. Felleisen, M., Finler, R.B., Flatt, M.: Semantics Engineering with PLT Redex. The MIT Press, Cambridge (2009)
6. Gabay, Y., Kfoury, A.J.: A calculus for java’s reference objects. SIGPLAN Not **42**(8), 9–17 (2007). <https://doi.org/10.1145/1294297.1294299>
7. Graham-Cumming, J.: CloudFlare’s new WAF: compiling to Lua. <https://blog.cloudflare.com/cloudflares-new-waf-compiling-to-lua> (2013). accessed 04 May 2020
8. Guha, A., Saftoiu, C., Krishnamurthi, S.: The essence of JavaScript. In: ECOOP ’10 (2010)
9. Haas, A., Rossberg, A., Schuff, D.L., Titzer, B.L., Holman, M., Gohman, D., Wagner, L., Zakai, A., Bastien, J.: Bringing the web up to speed with WebAssembly. In: Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, Association for Computing Machinery, New York, NY, USA, PLDI 2017, pp. 185–200. <https://doi.org/10.1145/3062341.3062363> (2017)
10. Ierusalimschy, R.: Programming in Lua. Lua.org (2003)
11. Ierusalimschy, R., de Figueiredo, L.H., Celes, W.: Lua—an extensible extension language. Software **26**(6), 635–652 (1996)

¹⁶ <https://sourcegraph.com/search?q=context:global+lang:Lua+%22+goto+%22+count:1000000&patternType=regexp&case=yes>.

12. Ierusalimsky, R., de Figueiredo, L.H., Celes, W.: The evolution of an extension language: a history of lua. In: Brazilian Symposium on Programming Languages (2001)
13. Ierusalimsky, R., de Figueiredo, L.H., Celes, W.: Lua 5.2 Reference Manual. www.lua.org/manual/5.2/manual.html (2013)
14. Igarashi, A., Pierce, B.C., Wadler, P.: Featherweight Java: a minimal core calculus for Java and GJ. *TOPLAS* **23**, 396–450 (2001)
15. Klein, C.: Randomized testing in PLT Redex. In: Proc. Scheme and Functional Programming, pp. 26–36 (2009)
16. Klein, C., McCarthy, J., Jaconette, S., Fidler, R.B.: A semantics for context-sensitive reduction semantics. In: APLAS'11 (2011)
17. Klein, C., Clements, J., Dimoulas, C., Eastlund, C., Felleisen, M., Flatt, M., McCarthy, J.A., Raskind, J., Tobin-Hochstadt, S., Fidler, R.B.: Run your research: On the effectiveness of lightweight mechanization. In: Proceedings of the 39th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, ACM, New York, NY, USA, POPL '12, pp 285–296, <https://doi.org/10.1145/2103656.2103691> (2012)
18. Krebbers, R., Wiedijk, F.: Separation logic for non-local control flow and block scope variables. In: FOSSACS'13, https://doi.org/10.1007/978-3-642-37075-5_17 (2013)
19. Leal, M.A., Ierusalimsky, R.: A formal semantics for finalizers. *J UCS* **11**(7), 1198–1214 (2005). <https://doi.org/10.3217/jucs-011-07-1198>
20. Lin, H.: Operational Semantics for Featherweight Lua. Master's thesis, San José State University (2015)
21. Lua Dev Team: Lua 5.2 test suite. <https://www.lua.org/tests/> (2013). Accessed 04 May 2020
22. Lua Dev Team: Lua 5.2 reference manual. <https://www.lua.org/manual/5.2/manual.html> (2015). Accessed 04 May 2020
23. Lua Developers: Expression as statements. <http://lua-users.org/wiki/ExpressionsAsStatements> (2009). Accessed 04 May 2020
24. Lua Developers: Lua analyzers. <http://lua-users.org/wiki/ProgramAnalysis> (2014). Accessed 04 May 2020
25. Lua Developers: Lua implementations. <http://lua-users.org/wiki/LuaImplementations> (2018). Accessed 04 May 2020
26. Lua Developers: Lua directory. <http://lua-users.org/wiki/LuaDirectory> (2020). Accessed 04 May 2020
27. Lua Developers: Uses. <https://www.lua.org/uses.html> (2017). Accessed 04 May 2020
28. LuaTex: Luatex. <http://www.luatex.org/languages.html> (2018). Accessed 04 May 2020
29. Maffei, S., Mitchell, J.C., Taly, A.: An operational semantics for JavaScript. In: APLAS '08 (2008)
30. Manura, D.: Vararg the second class citizen. <http://lua-users.org/wiki/VarargTheSecondClassCitizen> (2007). Accessed 04 May 2020
31. Mascarenhas de Queiroz, F.: Optimized compilation of a dynamic language to a managed runtime environment. PhD thesis, Pontifícia Universidade Católica do Rio de Janeiro (2009)
32. Moura, A., Rodriguez, N., Ierusalimsky, R.: Coroutines in lua. *J. Univers. Comput. Sci.* **10**(7), 910–25 (2004)
33. Pierce, B.C.: Types and Programming Languages, 1st edn. The MIT Press, Cambridge (2002)
34. Politz, J.G., Carroll, M.J., Lerner, B.S., Pombrio, J., Krishnamurthi, S.: A tested semantics for getters, setters, and eval in JavaScript. In: DLS '12 (2012)
35. Politz, J.G., Martinez, A., Milano, M., Warren, S., Patterson, D., Li, J., Chitipothu, A., Krishnamurthi, S.: Python: the full monty: a tested semantics for the Python programming language. In: OOPSLA '13 (2013)
36. Rossberg, A., (eds): Webassembly specification. <https://webassembly.github.io/spec/core/> (2021). Accessed 20 Feb 2021
37. Soldevila, M., Ziliani, B., Silvestre, B., Fridlender, D., Mascarenhas, F.: Decoding Lua: Formal semantics for the developer and the semanticist. In: Proceedings of the 13th ACM SIGPLAN Dynamic Languages Symposium, DLS 2017 (2017)
38. Soldevila, M., Ziliani, B., Fridlender, D.: Understanding Lua's garbage collection: Towards a formalized static analyzer. In: Proceedings of the 22nd International Symposium on Principles and Practice of Declarative Programming, PPDP 2020 (2020)

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.