Position Paper: Some Goals of the Luau Type System

ANDY FRIESEN, ALAN JEFFREY, and OTHER PEOPLE?, Roblox, USA

Luau is the scripting language used in creating Roblox experiences. It is a typed language based on the dynamically-typed Lua language, and uses type inference to infer types. These types are then in the IDE, for example when providing autocomplete suggestions. In this paper, we describe some of the goals of the Luau type system, focusing on where the goals are different from those of other type systems.

1 INTRODUCTION

The Roblox [7] platform allows anyone to create shared, immersive, 3D experiences. At the time of writing, there are approximately 20 million experiences available on Roblox, created by 8 million developers. Roblox creators are often young, for example there are over 200 Roblox kids' coding camps in 65 countries listed at [6].

The Luau programming language [5] is the scripting language used by developers of Roblox experiences. Luau is derived from the Lua programming language [2], with additional capabilities, including a type inference engine.

This paper will discuss some of the goals of the Luau type system, focusing on where the goals are different from those of other type systems.

2 HUMAN ASPECTS

2.1 Heterogenous developer community

Quoting a Roblox 2020 report [4]:

- Adopt Me! now has over 10 billion plays and surpassed 1.6 million concurrent users in game earlier this year.
- Piggy, launched in January 2020, has close to 5 billion visits in just over six months.
- There are now 345,000 developers on the platform who are monetizing their games.

This demonstrates how heterogenous the Roblox developer community is: developers of experiences with plays measured in billions are on the same platform as children first learning to code. Moreover, *both of these groups are important*, as the professional development studios bring high-quality experiences to the platform, and the beginning creators contribute to the energetic creative community.

Fig. 1. Roblox Studio's 3D environment editor (a), and script editor (b)

2.2 Goal-driven learning

All developers are goal-driven, but this is especially true for learners. A learner will download Roblox Studio (the IDE) with an experience in mind, often designing an obstacle course (an "obby") to play in with their friends.

The user experience of developing a Roblox experience is primarily a 3D interactive one, seen in Fig. 1(a). The user designs and deploys 3D assets such as terrain, parts and joints, and provides them with physics attributes such as mass and orientation. The user can interact with the experience in Studio, and deploy it to a Roblox server so anyone with the Roblox app can play it.

At some point during experience design, the user of Studio has a need which can't be met by the physics engine alone. "The stairs should light up when a player walks on them" or "a firework is set off every few seconds." At this point they will discover the script editor, seen in Fig. 1(b), and the Luau programming language.

This onboarding experience is different from many initial exposures to programming, in that by the time the user first opens the script editor, they have already built much of their creation, and have a very specific concrete aim. It suggests a Luau goal for helping the majority of creators: *support learning how to perform specific tasks* (for example through autocomplete suggestions and documentation).

2.3 Type-driven development

Professional development studios are also goal-directed (though the goals may be more abstract, such as "decrease user churn" or "improve frame rate") but have needs that are less common in learners:

- *Code planning*: code spends much of its development time in an incomplete state, with holes that will be filled in later.
- *Code refactoring*: experiences evolve over time, and it easy for changes to break previously-held invariants.
- *Defect detection*: code has errors, and detecting these at runtime (for example by crash telemetry) can be expensive and recovery can be time-consuming.

Detecting defects ahead-of-time is a traditional goal of type systems, resulting in an array of techniques for establishing safety results, surveyed for example in [3]. Supporting code planning and refactoring are some of the goals of *type-driven development* [1] under the slogan "type, define,

refine". For example. a common use of type-driven development is to rename a property, which is achieved by changing the name in one place, and then fixing the resulting type errors—once the type system stops reporting errors, the refactoring is complete.

To help support the transition from novice to experienced developer, types are introduced gradually, through API documentation and type discovery. Type inference provides many of the benefits of type-driven development even to creators who are not explicitly providing types.

3 TYPES

3.1 Infallible types

Goal: *support type-driven tools in all programs*.

Programs spend much of their time under development in an incomplete state, even if the final arifact is well-typed. Type-driven tools should support this, by providing type information for all programs. An analogy is infallible parsers, which perform error recovery and provide an AST for all input texts.

Program analysis can still flag type errors, for example with red squiggly underlining. Formalizing this, rather than a judgement $\Gamma \vdash M : T$, for an input terms M, there is a judgement $\Gamma \vdash M \Rightarrow M' : T$ where M' is an output term where some subterms are flagged $\underline{M}$. Write $\mathrm{erase}(M)$ for the result of erasing flagged type errors: $\mathrm{erase}(\underline{M}) = \mathrm{erase}(M)$.

The goal of infallible types is that every term can be typed:

- *Typability*: for every M and Γ, there are M' and T such that $\Gamma \vdash M \Rightarrow M' : T$.
- *Erasure*: if $\Gamma \vdash M \Rightarrow M' : T$ then $\mathrm{erase}(M) = \mathrm{erase}(M')$

Some issues raised by infallible types:

- Which heuristics should be used to provide types for flagged programs? For example, could one use minimal edit distance to correct for spelling mistakes in field names?
- How can we avoid cascading type errors, where a developer is faced with type errors that are artifacts of the heuristics rather than genuine errors?
- How can the goals of an infallible type system be formalized?

Related work: lots on type error reporting [?], and on heuristics for program repair [?], but not type repair, or on the semantics of programs with type errors.

3.2 Strict types

Goal: *no false negatives*.

For developers who are interested in defect detection, Luau provides a *strict mode*, which acts much like a traditional, sound, type system. This has the goal of "no false negatives" that is any run-time error is flagged.

The usual presentation of type safety is using type preservation and progress. This requires:

- *Operational semantics*: a reduction judgement $M \rightarrow N$ on terms.
- *Values*: a subset of terms representing a successfully completed evaluation.

We then represent error states as stuck states (terms that are not values but cannot reduce), and showing that no well-typed program is stuck. This is not true if typing is infallible, but can fairly straightforwardly be adapted:

- *Progress*: if $\vdash M \Rightarrow M'$, then either $M \rightarrow N$ or M is a value or M' is flagged.
- *Preservation*: if $\vdash M \Rightarrow M'$ and $M \rightarrow N$ and $\vdash N \Rightarrow N'$ and N' is flagged, then M' is flagged.

Some issues raised by infallible types:

- How should the judgements and their metatheory be set up?
- How should generic functions be handled?

- What does type inference or bidirectional typechecking look like in this setting?

Related work: blame analysis [?].

3.3 Nonstrict types

Goal: *no false positives.*

For developers who are not interested in defect detection, type-driven tools such as autocomplete can still be useful, and type-directed development can still be useful. For such developers, Luau provides a *nonstict mode*, which we hope will eventually be useful for all developers. This does *not* aim for soundness, but instead has the goal of "no false positives", in the sense that any flagged code is guaranteed to produce a runtime error when executed.

On the face of it, this is undecidable, since a program such as (if $f()$ then error end) will produce a runtime error when $f()$ is true, but we can aim for a weaker property, thjat all flagged code is either dead code or will produce an error. Either of these is a defect, so deserves flagging, even if the tool does not know which. For example, using this definition (if $f()$ then <u>error</u> end) is not a false positive, but (if <u>$f()$ then error end</u>) might be.

We can formalize this by defining an *evaluation context* $\mathcal{E}[\bullet]$, defining a *redex* of M to be a subterm N such that $M = \mathcal{E}[N]$, and defining $M\Uparrow$ when there is no value V such that $M \rightarrow^* V$. We can then define:

- *Snappy name*: if $\vdash M \Rightarrow M'$ and M' has a flagged redex then $M\Uparrow$.

Some issues raised by nonstrict types:

- Under this definition, any function that will terminate is unflagged, so flagging will often move from function definitions to call sites.
- This definition will not allow an unchecked use of an optional value to be flagged, for example if we define function $g()$ if $f()$ then return 5 else return nil end then a strict type system can flag $1 + g()$ but a nonstrict one cannot.
- Property update of tables in languages like Luau always succeeds (the property is inserted if it did not exist), which interacts badly with compositional analysis. For example, if module A exports a table $t = \{p = 5\}$, module B assumes $t.p$ is a number, and module C sets $t.p = $ false, then we cannot flag C since the assignment will succeed, and we cannot flag A or B without some whole-program analysis.
- The natural formulation of function types in a nonstrict setting is that of [?]: if $f : T \rightarrow U$ and $f(v) \rightarrow^* w$ then $v : T$ and $w : U$. This formulation is *covariant* in T, not *contavariant*.

Related work: success types [?] and incorrectness logic [?].

3.4 Mixing types

Goal: *support mixed strict/nonstrict development.*

Like every active software community, Roblox developers share code with one another constantly. First- and third-party developers alike frequently share entire software packages written in Luau. To add to this, many Roblox games are authored not by just one developer, but a whole team.

It is therefore crucial that we offer first-class support for mixing code written in strict and nonstrict modes.

Some issues raised by mixed-mode types:

- How much feedback can we offer for a nonstrict script that is importing strict-mode code?
- In strict mode, how do we talk about values and types that are drawn from nonstrict code?
- How can we combine the goals of strict and nonstrict types?
- Can we have strict and non-strict mode infer the same types, only with different flagging?

Related work: none???

4 CONCLUSIONS

In this paper, we have presented some of the goals of the Luau type system, and how they map to the needs of the Roblox creator community. We have sketched what a solution might look like, all that remains is to draw the damn owl [?].

REFERENCES

[1] Edwin Brady. 2017. *Type-Driven Development with Idris*. Manning.
[2] Lua.org and PUC-Rio. 2021. The Lua Programming Language. https://lua.org
[3] Benjamin C. Pierce. 2002. *Types and Programming Languages*. MIT Press.
[4] Roblox. 2020. Roblox Developers Expected to Earn Over $250 Million in 2020; Platform Now Has Over 150 Million Monthly Active Users. https://corp.roblox.com/2020/07/roblox-developers-expected-earn-250-million-2020-platform-now-150-million-monthly-active-users/
[5] Roblox. 2021. The Luau Programming Language. https://luau-lang.org
[6] Roblox. 2021. Roblox Education: All Educators. https://education.roblox.com/en-us/educators
[7] Roblox. 2021. What is Roblox. https://corp.roblox.com