



Automated test generation for SCRATCH programs

Adina Deiner¹ · Patric Feldmeier¹ · Gordon Fraser¹  · Sebastian Schweikl¹ · Wengran Wang²

Accepted: 20 October 2022 / Published online: 13 May 2023
© The Author(s) 2023

Abstract

The importance of programming education has led to dedicated educational programming environments, where users visually arrange block-based programming constructs that typically control graphical, interactive game-like programs. The SCRATCH programming environment is particularly popular, with more than 90 million registered users at the time of this writing. While the block-based nature of SCRATCH helps learners by preventing syntactical mistakes, there nevertheless remains a need to provide feedback and support in order to implement desired functionality. To support individual learning and classroom settings, this feedback and support should ideally be provided in an automated fashion, which requires tests to enable dynamic program analysis. In prior work we introduced WHISKER, a framework that enables automated testing of SCRATCH programs. However, creating these automated tests for SCRATCH programs is challenging. In this paper, we therefore investigate how to automatically generate WHISKER tests. Generating tests for SCRATCH raises important challenges: First, game-like programs are typically randomised, leading to flaky tests. Second, SCRATCH programs usually consist of animations and interactions with long delays, inhibiting the application of classical test generation approaches. Thus, the new application domain raises the question of which test generation technique is best suited to produce high coverage tests capable of detecting faulty behaviour. We investigate these questions using an extension of the WHISKER test framework for automated test generation. Evaluation on common programming exercises, a random sample of 1000 SCRATCH user programs, and the 1000 most popular SCRATCH programs demonstrates that our approach enables WHISKER to reliably accelerate test executions, and even though many SCRATCH programs are small and easy to cover, there are many unique challenges for which advanced search-based test generation using many-objective algorithms is needed in order to achieve high coverage.

Keywords Search-based testing · Block-based programming · SCRATCH

Communicated by: Dan Hao

Authors listed in alphabetical order.

✉ Gordon Fraser
Gordon.Fraser@Uni-Passau.De

¹ University of Passau, Innstr. 33, Passau, 94032, Germany

² North Carolina State University, 890 Oval Dr, Raleigh, NC 27606, USA

1 Introduction

Computer programming nowadays is at the forefront of education: Not only is programming considered an important skill that is included in general computer science education, it also plays a central role in the teaching of computational thinking (Lee et al. 2011). The term “computational thinking” refers to the ability to think or solve problems based on computing methods, and includes aspects such as abstraction, data representation, and logically organising data. A core vehicle to teach these aspects is programming. Since computational thinking is increasingly integrated into core curricula at primary school level, even the youngest learners nowadays learn how to create simple computer programs.

Teaching young learners programming requires dedicated programming languages and programming environments. Common novice programming environments, such as SCRATCH (Maloney et al. 2010), SNAP (Harvey et al. 2013), and ALICE (Cooper et al. 2000) engage young learners by allowing them to build programming artefacts such as apps and games, which connects computation with their real-world interests (Papert 1980). Novice programming environments typically have two distinguishing features: First, to avoid the necessity to memorize and type textual programming commands as well as the common syntactic overhead caused by braces or indentation, programs are created visually by dragging and dropping block-shaped commands from “drawers” containing all possible blocks. The blocks have specific shapes and only matching blocks snap together, such that it is only possible to produce syntactically valid programs. Second, the programs typically control graphical sprites in a game-like, interactive environment. Accordingly, many programming commands are high-level statements that control the behaviour of these graphical sprites.

While these simplifications and application scenarios reduce complexity and make programming accessible and engaging, the learning process is nevertheless challenging from multiple points of view: Learners may struggle implementing programs due to misconceptions (Sirkiä and Sorva 2012), and even though there are no syntax errors there still is an infinite number of possible ways to assemble blocks in incorrect ways (Frädrich et al. 2020). Teachers therefore need to support their students, but to do so they need to comprehend each individual learner’s program, which can be a daunting task in the light of large classrooms. Consequently, there is a need to support learners and teachers with the help of automated tools.

A primary means to enable automated tools to inform learners and teachers is by testing the programs. Given insights on the implemented behaviour, automated tools can identify missing or incorrect functionality, they can suggest which parts of the program to fix, how to fix them, or which steps to perform next in order to solve the overall task. A common prerequisite, however, is that tests can be automated. In the context of SCRATCH programs, the WHISKER framework (Stahlbauer et al. 2019) provides a means to automate testing. A WHISKER test automatically sends user events such as key presses or mouse clicks to the program, and observes the resulting behaviour. However, creating WHISKER tests is a challenging task. For example, the test suite for a simple fruit catching game used in the original WHISKER study (shown in Fig. 1) consists of 869 lines of JavaScript code. Some of the complexity of creating such tests can be alleviated by providing suitable user interfaces and more abstract means to specify the tests (Wang et al. 2021b). However, at the end of the day the tests nevertheless require non-trivial manual labour, which is problematic considering the likely target audience of teachers who may not be adequately educated software engineers.

In this paper we aim to mitigate this problem by relieving users of the task of creating tests themselves. Given a SCRATCH program, we aim to automatically generate a set of WHISKER tests that execute all parts of the code. The program under test could be an example solution for a given task, such that the test suite can then be executed against all student solutions. Alternatively, the program could be a student solution for which a dynamic analysis is desired. However, even though SCRATCH programs tend to be small and playful, generating tests for them automatically is nevertheless challenging. In an initial proof of concept, we demonstrated the feasibility of using search techniques to automatically generate sequences of interactions with SCRATCH programs (Deiner et al. 2020), but also revealed multiple obstacles that make test generation difficult: Unlike traditional code, SCRATCH program executions tend to take substantial time due to the frequent use of motion- and sound-related animations and time-delays. Automated test generation techniques, however, rely on frequently executing programs. The original WHISKER test execution framework (Stahlbauer et al. 2019) reduced non-determinism by controlling random number generators, but experience has shown that the timing aspects of these programs and the unpredictability of the scheduler when executing these highly concurrent programs nevertheless make deterministic executions difficult. Finally, there are technical challenges related to the questions of which interactions a program should receive, and which algorithm to use in order to explore possible sequences of such interactions.

In order to address these challenges, we extend our prior work on the WHISKER testing framework (Stahlbauer et al. 2019) and WHISKER test generation (Deiner et al. 2020). In detail, the contributions of this paper are as follows:

- We modify the execution model of the SCRATCH virtual machine that makes executions deterministic, even in the light of timing and concurrency. Given this modification, tests are fully reproducible and can be executed in a highly accelerated fashion.
- We propose a strategy that takes the source code as well as the runtime state of a program during its execution into account to determine which user events are suitable for interacting with a program under test.
- We adapt random and search-based test generation approaches to the scenario of generating high-coverage test suites for SCRATCH programs, which includes improvements to encoding, fitness function, and algorithms.
- We implement a full-fledged test generation framework that combines generated event sequences with regression assertions.
- We empirically study the correctness of our virtual machine modifications and the ability of our test generation techniques to cover and test the code of real SCRATCH programs using four datasets.




Our experiments demonstrate that, even when accelerating test execution by a factor of 10, test results are completely deterministic and devoid of any flakiness using our approach. While we find that many users create programs that are small and easy to cover fully, there are countless unique challenges to test generation for SCRATCH, ranging from extracting suitable events for exercising a program to calculating appropriate reachability estimates to guide test generation. Our extension of the WHISKER framework implements techniques that, collectively, allow many-objective search-based test generation to achieve an average of 95.5% coverage on common user-written programs, and on average 69.2% coverage on popular projects. Consequently, WHISKER represents an important step towards enabling dynamic analysis of SCRATCH programs and countless resulting possible applications in the area of supporting programming learners. To support researchers in

developing these applications, and to develop new techniques to improve coverage further, WHISKER and its mature automated test generation framework are freely available as open source software.

2 Background

2.1 SCRATCH Programs

SCRATCH programs revolve around a *stage* on which graphical *sprites* process user inputs and interact. Figure 1 shows a SCRATCH program with three sprites: bowl, bananas and apple; the stage contains the background image. Conceptually, we define a SCRATCH program as a set of *actors*, one of which is the stage and the others are sprites. Actors are rendered on a canvas; each actor is rendered on a separate layer (Stahlbauer et al. 2019). An actor is composed of sets of scripts, custom blocks, sound and image resources. The resources are used for example to provide background images on the stage, or to decorate sprites with costumes. The currently chosen costume is an example of an *attribute* of an actor, and other attributes include position, rotation, or size. Actors can also define *variables* which are untyped and contain numeric or textual data. Scripts consist of individual *blocks* stacked together. The SCRATCH language consists of different types of blocks with different shapes, and programs are arranged by combining blocks in ways that are permitted by their shapes.

- Hat blocks : Each script can have only one hat block, which represents an event handler that triggers the execution of the script. Scripts without hat blocks can only be executed by the user double clicking on them.
- Stack blocks , , ...: Regular program statements, for example to control the appearance or motion of sprites, can be stacked on top of each other. The stacking represents the order of the control flow between the statements.

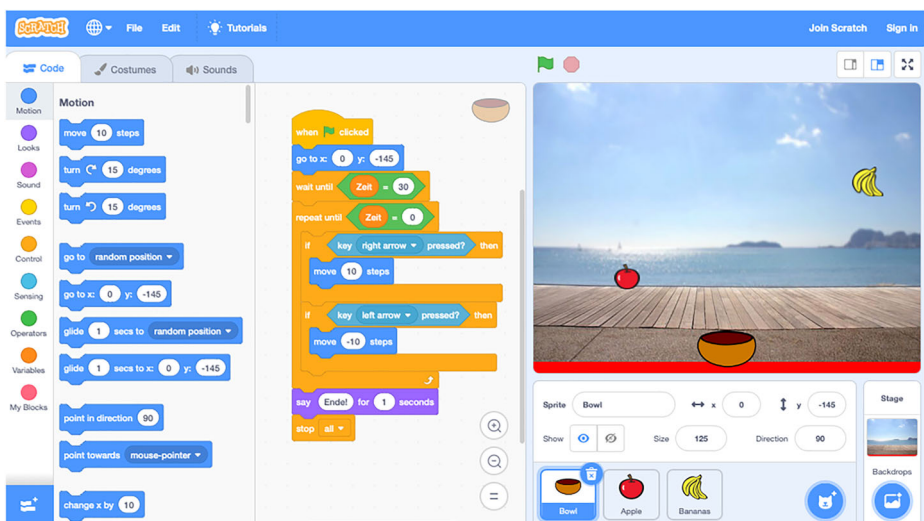












Fig. 1 The SCRATCH user interface

- C blocks : These blocks are named after their shape and represent control flow (if, if-else, loops). The conditionally executed code is contained within the C-shape.
- Reporter blocks : These blocks represent variables and expressions and can be used as parameters of other blocks.
- Boolean blocks : These are special reporter blocks that represent Boolean values.
- Cap blocks : These are blocks after which no stack blocks can be attached as they either terminate the execution or it never proceeds beyond them (e.g., forever loops).
- Custom blocks  : These blocks are essentially macro scripts. An instance of a custom block triggers the execution of the corresponding macro script.

SCRATCH programs are executed in the SCRATCH virtual machine, and controlled by the user via mouse, keyboard, microphone, or other input devices. That is, a program can react to mouse movement, mouse button presses, keyboard key presses, sound levels, or entering answers to  blocks. In addition, there is a global Greenflag event  which represents the user starting the program through the green flag icon in the user interface (cf. Fig. 1). The SCRATCH language also contains broadcast statements , which trigger corresponding message receiver hat-blocks . The execution of a script is initiated when the event corresponding to its hat block occurs, resulting in a process p . Executing a SCRATCH program therefore results in the creation of a collection of concurrent processes P , and the state of each process is defined by the control location as well as the values of all variables and attributes of the actor.

Execution is operationalised by the *step* function of the virtual machine. Figure 2 shows a simplified version of one scheduling step performed by the SCRATCH VM to update its internal state. Each step has a predefined *step time* duration and starts by determining which scripts are currently active and have to be executed. The collection of active scripts $P' \subset P$ consists of processes triggered by recent user inputs and already active scripts from previous time steps. All active processes are then handed over for execution to the *sequencer*. The sequencer mimics parallelism by sequentially executing all received processes in batches

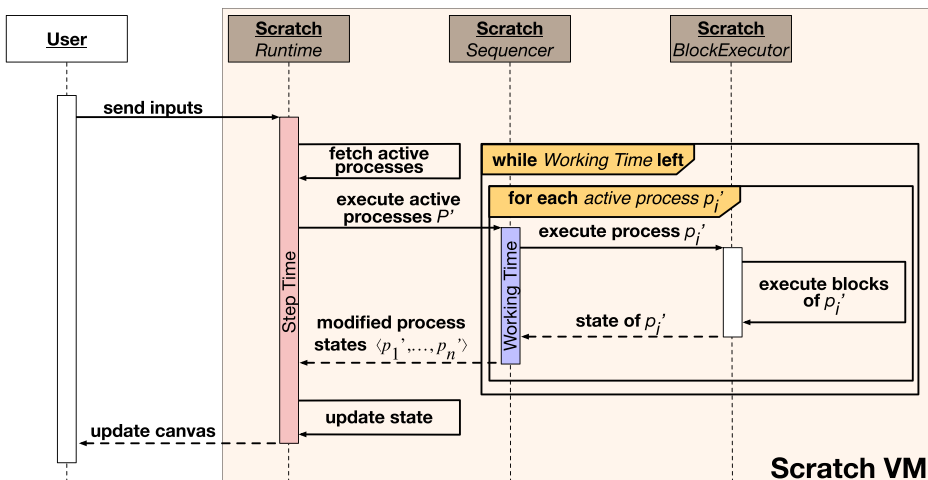












Fig. 2 Simplified scheduling function of the SCRATCH VM

$B = \langle p'_1, \dots, p'_n \rangle$ until the *working time*, which is by default set to two-thirds of the *step time*, has elapsed. In order to avoid non-deterministic behaviour, the execution of a process batch is never interrupted, even if the *working time* has depleted. Whenever a single process p' is scheduled for execution, the process is transferred to the *block executor*. Upon receiving a script to execute, the block executor processes each block of the given process p' until all blocks of the script's process have been executed or specific blocks, forcing the process to halt, are encountered. These process halting blocks consist of:

-  and  blocks, which force the process to wait until a user-defined timeout x has run out or some condition cond is met.
-  blocks, which create a think/speech bubble for the specified amount of time x on top of the sprite containing the block.
-  blocks, which move the given sprite gradually within a time frame x to a specified location.
-  blocks, which force the program execution to halt until the defined sound file x has been played completely.
-  blocks, which work like a  block by translating the given text argument x into the sound file y .
- The last block contained within ,  and  blocks, which forces the process to halt until the next process batch is executed.

Eventually, the state of the currently executed process p' is reported back to the sequencer. As soon as the *working time* has depleted and the full batch of processes has been executed, the collection of modified process states $\langle p'_1, \dots, p'_n \rangle \in P'$ is handed back to the runtime environment. Finally, the runtime environment updates the internal state of the SCRATCH VM and notifies the user by redrawing the canvas to meet the state changes.

2.2 The WHISKER Testing Framework

Testing a program means executing the program, observing the program's behaviour, and checking this behaviour against expectations. WHISKER (Stahlbauer et al. 2019) automates this process for SCRATCH programs: Conceptually, a WHISKER test consists of a *test harness*, which sends user events to the SCRATCH program under test, and a set of SCRATCH *observers*, which encode properties ϕ that should be checked on the program under test.

As illustrated by Fig. 3, WHISKER executes tests by wrapping the SCRATCH VM's scheduling function and inheriting its *step time*. First, WHISKER queries the test harness for an input to be sent to the SCRATCH program under test. Then it performs a *step* by sending the obtained input in the form of an event to the SCRATCH program. The SCRATCH VM then invokes its scheduling function (Fig. 2). After the *working time* has expired, the scheduling function stops and reports the new state back to WHISKER. This state is handed over to the test observer, which checks if the actual state matches the expected properties ϕ . If it does not, then a failing has been found and is reported in the form of an *error witness* (Diner et al. 2021), which contains the whole input sequence leading to the violation of the given property.

A static test harness provides inputs encoded in JavaScript to the program. Arbitrary events can be sent based on time intervals or when certain conditions hold. As an example,

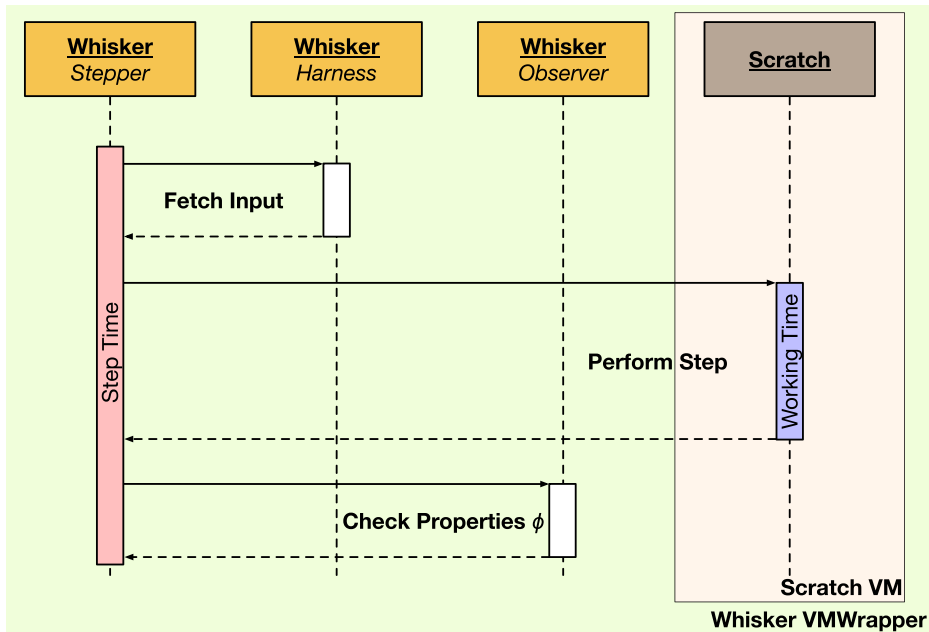


Fig. 3 Execution step of the WHISKER VMWrapper: Select an input, send that input to the SCRATCH VM, and finally match the resulting state against the expected properties ϕ

Listing 1 shows a WHISKER test for the project in Fig. 1. The test consists of pressing the left cursor key for ten steps and checking whether the bowl-sprite has moved to the left.

WHISKER also supports dynamic test harnesses, where the program is exercised with randomly generated sequences of inputs. Although these are often sufficient to fully cover simple programs, previous work (Stahlbauer et al. 2019) has shown that more complex programs are not always fully covered. In addition, in cases where an instructor or a researcher needs to author tests for multiple students' solutions in one assignment, defining one set of inputs is almost never sufficient to cover all different student programs, where each student may have their own unique implementation for properties such as actor movement speed and game begin/end conditions (Wang et al. 2021a). Finally, specifying SCRATCH observers may be easier for a given sequence of user inputs in contrast to specifying the expected outcome for arbitrary sequences inputs. Therefore, the aim of this paper is to generate static test harnesses, i.e., test suites that reach all statements of a program under test.

```










const test = async function(t) {
  let sprite = t.getSprite('Bowl');
  let oldX = sprite.x;
  t.keyPress('Left', 10);
  await t.runForSteps(10);
  t.assert.isTrue(sprite.x < oldX);
}
  
```

Listing 1 Example WHISKER test case for the game shown in Fig. 1

3 Accelerated and Deterministic Test Execution

The nature of SCRATCH programs causes two issues for automated testing: First, the frequent use of animations and timed behaviour causes executions to take long. Second, this time-dependent behaviour, the randomised nature of games, and various implementation aspects of SCRATCH may lead to non-determinism. To enable efficient and reliable testing of SCRATCH programs we therefore modified the SCRATCH VM to decrease the execution time and to make executions deterministic. Additionally, since some blocks check for sound originating from a device's microphone, we modified the SCRATCH VM to allow WHISKER to generate virtual sound levels without requiring a real physical microphone.

3.1 Accelerating Execution

To increase test execution speed we apply two essential modifications to the SCRATCH VM: First, the rate at which inputs are sent to the SCRATCH VM is increased; second, all blocks that halt the execution of a process to wait for time to pass (see Section 2.1) are instrumented to reduce their time-dependent arguments in proportion to the chosen acceleration factor. While the SCRATCH VM tries to execute process batches until the working time has elapsed, due to the nature of SCRATCH programs most processes sooner or later hit some statement that causes waiting. Indeed, we have observed that processes tend to spend more time waiting than executing within a step. Therefore, to effectively increase the rate at which inputs are sent to the SCRATCH VM, we reduce the step time of the WHISKER VMWrapper by the selected acceleration factor. Since the step time is directly linked to the SCRATCH VM's working time, as explained in Section 2.1, reducing the step time automatically leads to an increased update interval of the SCRATCH VM. Thus, a decrease in the step time results in sending inputs more frequently to the program under test and furthermore increases the rate at which the SCRATCH VM processes these inputs. To illustrate the increased execution speed, Fig. 4 depicts an accelerated version of WHISKER's scheduling function shown in Fig. 3. In the accelerated scheduling function, an acceleration factor of two cuts the step and working time in half, resulting in twice as many steps within the same time frame. However, some blocks force a process to enter an execution halting process state until a timeout or some user-defined condition is met (Section 2.1). Hence, even if the speed at which the SCRATCH VM updates its internal state is increased, these blocks would still force the process they are contained in to halt for the given amount of time. Therefore, statements that contain a time-dependent argument such as , ,  and  are instrumented to reduce their time argument x by the given acceleration factor. On the other hand, execution halting blocks that do not directly hold any time-dependent arguments like  and  are accelerated by reducing the play duration of the given or translated sound file appropriately. Lastly, statements that stop the execution of a process until a specific condition  is met, as the  block, are not altered at all because these conditions  already emerge earlier due to accelerated program execution. By applying both VM modifications to the original SCRATCH VM, we construct a modified SCRATCH VM used by WHISKER, capable of effectively increasing test execution speed by a user-defined acceleration factor.

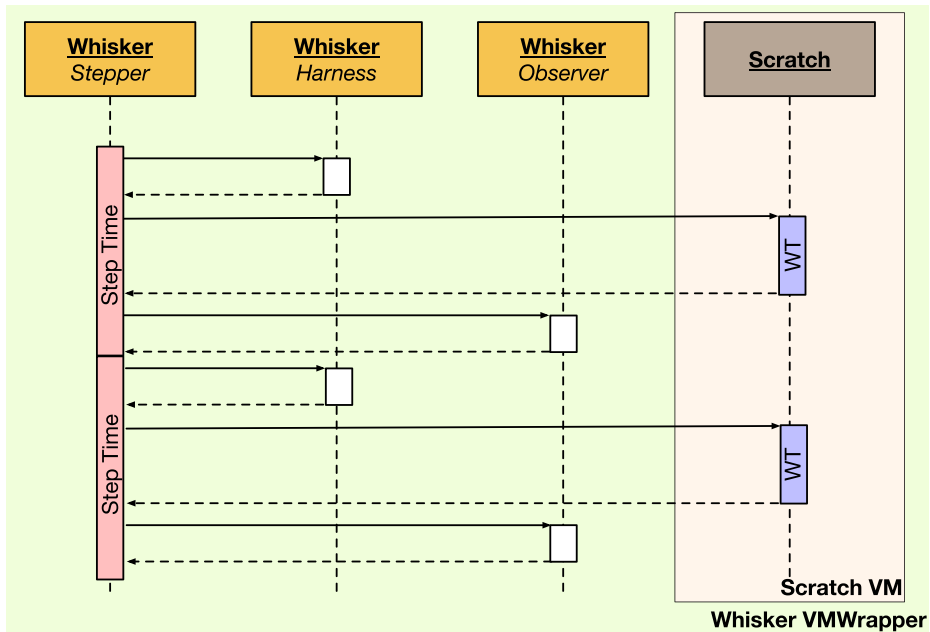


Fig. 4 Accelerating a SCRATCH program by halving WHISKER's step time and the SCRATCH VM's working time ("WT")

3.2 Ensuring Determinism

Most of the projects created within the SCRATCH programming environment represent simple games. A very prominent characteristic of games is their use of random number generators, which results in non-deterministic behaviour. Unfortunately, this frequently leads to flaky test suites (Luo et al. 2014; Gruber et al. 2021), which is undesirable in any testing tool but especially problematic in WHISKER's application scenario: For example, a WHISKER test related to a graded assignment could pass on a student's system but fail on the teacher's machine, leading to a potentially unfair grading process. In order to avoid flaky tests originating from randomised program behaviour, WHISKER offers the option to seed the SCRATCH VM's random number generator to a user-defined seed, by replacing the global `Math.random` function with a seeded one at runtime. However, random number generators do not pose the only source of non-deterministic behaviour. In consequence of its many temporal dependencies, the SCRATCH VM itself is also susceptible to non-determinism. This becomes especially apparent in programs containing blocks that take temporal values as an argument. To exemplify this problem, consider Fig. 5, showcasing a program consisting of an elephant (originating from a study by Geldreich et al. (2016)) that changes its costume (and thus its visual appearance) every second. When executing this program for the same amount of time on different machines, the last selected costume might not always be the same.

This non-deterministic behaviour originates from the way the sequencer repeatedly executes active processes. As depicted in Fig. 2, the sequencer obtains all currently active processes P' from the runtime environment and keeps executing batches of those processes until the allocated *working time* has elapsed. The root cause of non-deterministic behaviour

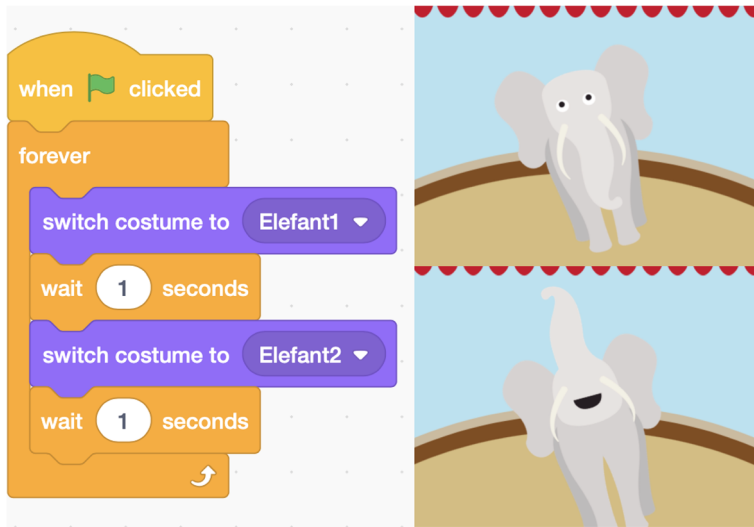




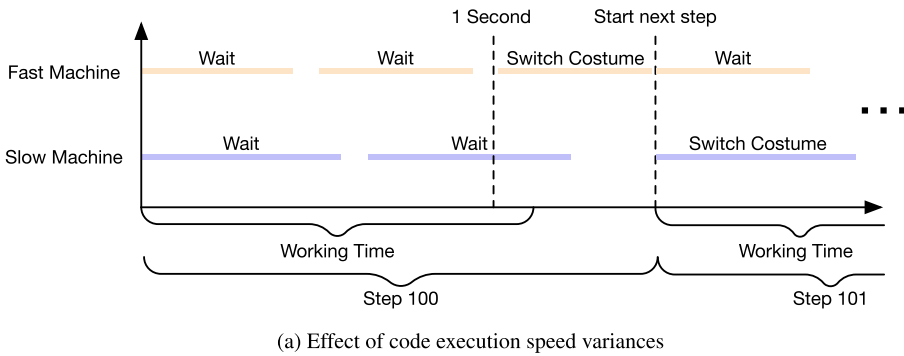
Fig. 5 SCRATCH project containing an elephant that changes its visual appearance every second

resides in the sequencer processing each batch of processes $b \in B$ as a whole and its inability to interrupt the execution of a batch b even if the allocated *working time* has run out. Therefore, differences in code execution speed between systems of diverging performances will allow the sequencer to step through and execute a varying number of process batches in the course of one *working time* interval. Hence, fast machines have a higher chance of striking a specific time interval earlier than slow machines. As depicted in Fig. 6, for the elephant example, this behaviour eventually leads to earlier costume changes in faster machines.


Accelerating execution may facilitate non-deterministic behaviour since the window of acceptable execution speed variances shrinks proportionally to the acceleration factor used. For example, an acceleration factor of 5 reduces the wait duration in the elephant project from one second down to 0.2 seconds. With a wait duration of only 0.2 seconds, the probability of more performant machines repeatedly triggering a state change faster than less performant ones increases. Moreover, as illustrated in Fig. 4, higher acceleration factors decrease the *working time* interval, further enforcing diverging program behaviour.

To eliminate non-deterministic behaviour originating from the scheduling function, we further modify the accelerated SCRATCH VM established within Section 3.1: First, time-dependent arguments within specific SCRATCH blocks are replaced with a discrete measure which is added to the SCRATCH VM and based on the number of executed steps so far. Since the SCRATCH VM implements different ways to treat time in different types of blocks, there are multiple different ways this change has to be implemented. Second, in a similar way, time-dependent WHISKER event parameters, defining how long inputs should be sent to the SCRATCH VM, are modified to be based on the number of executed steps as well.

In order to replace the imprecise measurement of time, a step counter $sc \in \mathbb{N}$ is added to the SCRATCH VM's runtime environment. The step counter is decoupled from the exact unit of time measured in seconds and responsible for counting the number of steps WHISKER has executed so far. Whenever a time-dependent  block is encountered during the execution of a SCRATCH program, the temporal argument $x \in \mathbb{Q}_{\geq 0}$ of the waiting block measured in seconds is translated into the corresponding number of steps $s \in \mathbb{N}$. Converting



(b) State on fast machine










(c) State on slow machine


Fig. 6 Effect of code execution speed variances on the Elephant project with each beam representing one process batch $b \in B$: Fast machines execute process batches more frequently within one step and thus have a higher chance of striking the end of the 1 second block earlier than machines having a low execution speed, resulting in diverging elephant states




seconds into the appropriate number of steps is done by dividing the time with the fixed *step time*, which already incorporates the chosen acceleration factor. For example, concerning the `wait 1 seconds` blocks of the depicted elephant project and assuming a *step time* of 10ms, the duration of all waiting blocks x is transformed into $s = 1s/10ms = 100$ steps. At the time of entering the process halting block, the current step count sc_0 is added to the calculated number of steps s to obtain the step count value $sc_r = sc_0 + s$ at which execution of the halted process can resume. Every time the process containing the time-dependent block is executed again, the current step count sc is checked, and execution is resumed iff $sc > sc_r$.



Very similar to the wait block is the `timer` block, which measures the time spent since program execution started or the time passed since a `reset timer` block was encountered. Non-deterministic real-time measurements in timer blocks are replaced by a new variable, which is increased by a value of 0.075 after every executed step and set back to zero after encountering a reset block. The value of 0.075 was empirically derived by minimizing the difference to real-time measurements. Although those changes may introduce slight and, in most cases, non-perceivable behavioural differences to the original SCRATCH VM, they nonetheless ensure deterministic program behaviour and are therefore preferred over real-time measurements, which will always lead to flaky behaviour.

Whereas wait and timer blocks are realised in SCRATCH as custom timers, , ,  and  blocks are implemented differently and hence require special treatment.

Instead of maintaining a simple timer,  blocks withhold a *promise* until a timeout set via JavaScript's `setTimeout()` function has elapsed. The blocked process is kept in the Yield process state and can only resume its execution if it attains the withhold promise from the yield forcing block. In addition to the already present non-determinism caused by the SCRATCH VM, the `setTimeout()` function is known to be very inaccurate, therefore enforcing non-deterministic behaviour even more¹. Hence, it does not suffice to simply replace the timeout duration x with the corresponding number of steps s . To avoid the problematic JavaScript function, these blocks were modified to resemble a  block by setting the blocked process into the Wait process state instead of the Yield process state. Furthermore, in order to retain the functionality of  blocks, a speech/think bubble is placed before, and removed after, the simulated wait above the corresponding *actor*.

Blocks using  statements, on the other hand, repeatedly change an *actor*'s position on the canvas in relation to the elapsed time. These blocks are instrumented by transforming the total gliding duration x to the corresponding number of steps s . Additionally, whenever a specific glide block is entered for the first time, we calculate and store the glide terminating step count $sc_r = sc + s$, after which the respective block reaches its destination. Finally, by setting the the current step count sc in relation to the glide terminating step count sc_r , the position of the sprite can precisely be determined in each execution step.

Lastly,  and  blocks behave similar since both halt execution until the given or translated file x has been played. Because sound files specify their duration, the step count sc_r at which execution will resume can be calculated the same way as for the  blocks by translating the duration into the corresponding number of steps s .

Besides their translated sound file,  blocks contain another source of flaky behaviour by querying a remote server to produce a sound file for the text argument x . Due to network uncertainties (Luo et al. 2014), the translation of the text argument always takes different lengths of time, which in turn delays the calculation of sc_r and eventually leads to non-deterministic behaviour. To eliminate network uncertainties, we modified the SCRATCH VM to translate and cache all resulting sound files of  blocks during the loading process of the block hosting sprite, i.e., before the execution of the SCRATCH program.

By abstracting the measure of time x to logical execution steps s , we can exactly define at which step count sc_r a halted process can be resumed. As a consequence, locks originating from execution halting blocks can only be released between steps and no longer within a single step, which makes the *working time* obsolete. Therefore, the SCRATCH VM's *sequencer* is further modified to only execute a single process batch instead of executing as many process batches as possible within the *working time*. Since the implemented acceleration technique instruments execution halting blocks and because a single process is executed until the end of its corresponding script or until reaching a process halting block, the removal of the *working time* does not change the behaviour of the SCRATCH VM. Regarding the elephant project, as shown in Fig. 7, instead of changing the costume as soon as one second

¹<https://developer.mozilla.org/en-US/docs/Web/API/WindowOrWorkerGlobalScope/setTimeout>, last accessed September 2022

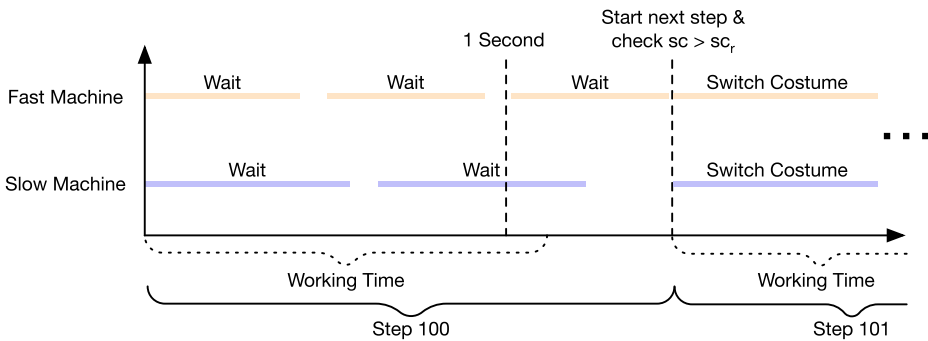




Fig. 7 Instead of waiting for exactly one second, the costume-changing blocks of the elephant now have to wait for exactly 100 steps until they are allowed to change their costume. The *working time* interval is no longer needed

has passed, the step counter ensures that the elephant will constantly change its costume every time precisely 100 steps have elapsed, no matter how fast the given machine executes process batches.

Besides the implementations of the SCRATCH blocks, the events that WHISKER tests consist of depend on temporal values as well, for example when deciding how long to send a key press to the program under test. To avoid diverging program executions, all temporal arguments x of WHISKER events are translated into steps s . Furthermore, to ensure all events can have impact on the SCRATCH VM, we enforce a minimal step duration of one step. Thus, given a step duration of 30ms and a `KeyPress` event with a duration of 29ms, the `KeyPress` event is translated into an event that presses the key for exactly $s = \frac{29}{30} \rightarrow 1$ step.

3.3 Virtual Sound

The hat block  as well as the sensing block  both check for the presence of a given sound level. When these blocks are executed, the SCRATCH VM tries to determine the current sound level in the user's environment by accessing the device's microphone. In case a microphone has been detected, the sound level is determined by calculating the *root mean square* (RMS) of the sound wave measured by the microphone. Then, the RMS value is scaled into the range $[0, 100]$, with 0 indicating that no noise has been detected and 100 representing the highest possible sound level.

In order to test SCRATCH projects without microphone (e.g., on a compute server), we added virtualised sound to the SCRATCH VM using a new variable *virtual sound*. The virtual sound variable is directly accessible via the WHISKER framework and thus can be utilised to simulate sound levels in the range of $[0, 100]$, mirroring the RMS value range. The sound event in WHISKER can send sound with a given volume for a predefined number of steps to the program under test. As soon as the defined number of steps have elapsed, WHISKER sets the virtual sound level to a value of -1 to indicate the end of the simulated sound sequence. To guarantee that a given program recognises the simulated sound, we further modified the SCRATCH VM to always check for the presence of virtual sound first before trying to access a microphone. However, to still allow the SCRATCH VM to fall back to its default behaviour using a physical microphone to detect sound, we use a value of -1 to indicate that no virtual sound is currently intended to be sent to the SCRATCH VM.

4 Test Generation for SCRATCH

A SCRATCH program processes streams of user inputs (e.g., keyboard and mouse events) to update program states. As an example, consider the SCRATCH program shown in Fig. 8, which contains two sprites, a *cat* and a *bear*. The program starts when a player clicks on the cat sprite. The cat first greets the bear by saying “Hello bear!” for 2 seconds. Afterwards, the script of the bear receives the message that they need to answer the cat, and then greets the cat back by saying “Hello cat!” for 2 seconds. Then, if the user presses the space key, the bear will change to a “smiling bear” costume. This program includes two scripts and two types of input events: clicking on the cat sprite and pressing the space key. Consequently, testing this program would involve repeatedly sending either of these two events to the program.

In order to systematically generate test inputs for SCRATCH programs, we derived a representative set E of input events and encoded them into WHISKER:

$$E = \{\text{Greenflag, KeyPress, ClickSprite, ClickStage, TypeText, TypeNumber, MouseDown, MouseMove, MouseMoveTo, DragSprite, Sound, Wait}\}$$

Some events require parameters, such as (x, y) coordinates for MouseMove. Thus, a user event is fully defined by a tuple (e, v) consisting of an event type $e \in E$ and a list of parameter values $v = \langle v_1, v_2, \dots, v_i \rangle$. Table 1 summarises the supported events and their parameters.

4.1 Test Generation

A test case $t = \langle e_1, e_2, e_3, \dots, e_i \rangle$ consists of a sequence of events and their parameters. In the simplest form of test generation, a test case may be produced by combining several randomly selected events from the sequence of available events E until the desired test case length has been reached. The randomTestGeneration Algorithm 1 repeatedly chooses an event from E and determines parameters for the chosen event using the random function, which randomly selects a single element from a set of elements S . Note we are excluding the Greenflag event from E since by design WHISKER automatically sends this event at the beginning of every test execution. While straight-forward to implement, this strategy might be ineffective for test generation as events can be chosen from E for which no corresponding handler exists in the project under test. For example, in Fig. 8, sending events of category MouseDown or Sound is pointless since the program cannot respond to these

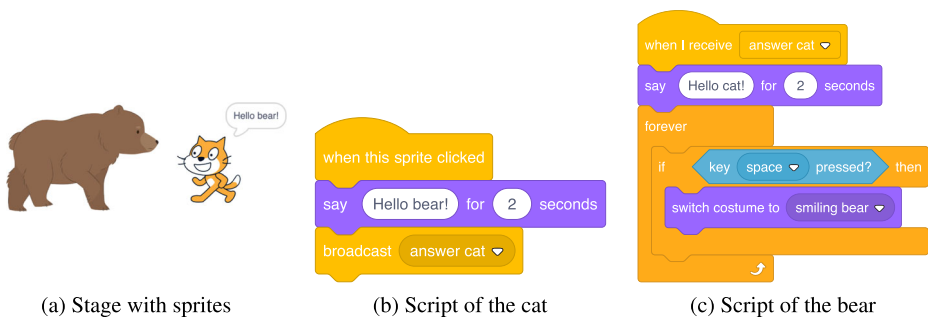


Fig. 8 Example SCRATCH program: The cat says “Hello bear!” when clicked and broadcasts the message “answer cat”. The bear receives this message and then says “Hello cat!”. Afterwards, when the space key is pressed the bear switches its costume to “smiling bear”

```
input  : number  $N \in \mathbb{N}$  of test cases to generate
output : test suite  $T = \langle t_1, \dots, t_N \rangle$  (a sequence of test cases)
function randomTestGeneration( $N$ ):
     $T \leftarrow \langle \rangle$ ;
    for  $i \leftarrow 1$  to  $N$  do
         $e \leftarrow \text{random}(E \setminus \langle \text{Greenflag} \rangle)$ ;
         $\langle p_1, \dots, p_n \rangle \leftarrow \text{parameter list of } e$ ;
         $v \leftarrow \langle \rangle$ ;
        for  $j \leftarrow 1$  to  $n$  do
             $V_j \leftarrow \text{value domain of parameter } p_j$ ;
             $v \leftarrow v \mathrel{++} \langle \text{random}(V_j) \rangle$ ;
         $T \leftarrow T \mathrel{++} \langle (e, v) \rangle$ ;
    return  $T$ ;
```

Algorithm 1 Random test generator

events. Thus, more fine-grained event extraction methods which filter irrelevant events from E are needed.

4.2 Event Extraction





The inclusion of irrelevant events can be avoided by considering only events for which event handlers exist in the source code. Event handlers can take on two different forms in SCRATCH. First, there are event handlers in terms of hat-blocks starting a script s (such as  and  in Fig. 8). Second, it is also possible to query the

Table 1 Events supported by WHISKER

Event	Parameter list	Description
Greenflag	—	Starts SCRATCH program execution.
KeyPress	Key and duration	Presses a certain key for a specified duration.
ClickSprite	Sprite	Clicks on a sprite or one of its clones.
ClickStage	—	Clicks on some location that is not occupied by sprites.
TypeText	Text	Enters a string from the pool of (1) all strings used in direct comparisons with the  block, (2) a randomly generated string, and (3) fixed seed strings ('0', '10', 'Hello').
TypeNumber	Number	Inserts the supplied number into a text field.
MouseDown	Status	Toggles the mouse button status (pressed/not pressed).
MouseMove	(x , y) coordinates	Moves the mouse to the target coordinates.
MouseMoveTo	Sprite	Moves the mouse to the position of the target sprite.
DragSprite	Sprite/colour/edge and angle	Drags the sensing sprite to a sensed sprite/colour or randomly selected edge on the canvas. Additionally, if angle ≤ 360 the determined dragging location is moved along the angle direction by the size of the sensing sprite.
Sound	Volume and duration	Sends virtualised sound to the SCRATCH VM for a specified duration
Wait	Duration	Waits for a specified duration.





















state of mouse and keyboard through sensing blocks. For example, in the program shown in Fig. 8, the sensing block  observes if the space key is pressed. Thus, a corresponding KeyPress event should be included in the sequence of available events E . Table 2 summarises the event handling blocks for all the supported user events.

Extracting only events for which corresponding event handlers exist in the SCRATCH program may still include irrelevant events: On the one hand, sensing blocks cannot process events if their scripts are inactive. For example, in Fig. 8 there is no point in sending KeyPress(space) events when the forever loop in script Fig. 8c is not executing. On the other hand, if an event which triggers a hat block of a currently active script is sent to the program, the execution of that script is halted and starts anew. Therefore, some sensing blocks potentially never become activated because their scripts are constantly restarted.

Consequently, `dynamicEventExtraction` considers the source code of the SCRATCH program as well as the current program state γ , as shown in Algorithm 2. The filtered events are formed by extracting events for which corresponding sensing blocks occur within active scripts and collecting events from hat blocks whose scripts are currently inactive. For example, in Fig. 8, at the beginning of the program execution, the event set only includes ClickSprite(cat). After 2 seconds, when the script in bear starts to execute, ClickSprite(cat) is removed from the event set, and KeyPress is added to the event set. This allows the input events to only include highly relevant events at each moment of program execution. Furthermore, because only active sensing blocks are considered, loose sensing blocks in unconnected scripts that have no hat block responsible for starting the execution of the script are automatically removed as well. Lastly, whenever an event $e \in \{\text{TypeText}, \text{TypeNumber}\}$ is extracted, we follow the lead of the SCRATCH environment, setting the focus of the user interface to the input field, and further restrict the sequence of available events to $\langle e, \text{Wait} \rangle$.

As the `dynamicEventExtraction` observes the current program state, it is also capable of

Table 2 Mapping of SCRATCH blocks to the corresponding WHISKER event handlers

Block	Event
 	KeyPress
	ClickSprite
	ClickStage
	TypeText
 followed by numeric operations on 	TypeNumber
	MouseDown
  	MouseMove
  	MouseMove
	MouseMoveTo
  	DragSprite
 	Sound

```

input   : the current program state  $\gamma$  from which events will be extracted
output  : set of events  $E_{dynamic}$ 
function dynamicEventExtraction( $\gamma$ ):
     $E_{dynamic} \leftarrow \langle \rangle$ ;
     $A \leftarrow$  currently active scripts in  $\gamma$ ;
    foreach script  $a \in A$  do
         $h \leftarrow$  hat block of  $a$ ;
        foreach block  $b \in a$  do
            if  $b \neq h$  then
                 $e \leftarrow$  extract event from  $b$  including inferable parameters;
                 $E_{dynamic} \leftarrow E_{dynamic} \uparrow\uparrow \langle e \rangle$ ;

     $I \leftarrow$  inactive scripts in  $\gamma$ ;
    foreach script  $i \in I$  do
         $h \leftarrow$  hat block of  $i$ ;
         $e \leftarrow$  extract event from  $h$  including inferable parameters;
         $E_{dynamic} \leftarrow E_{dynamic} \uparrow\uparrow \langle e \rangle$ ;

    if  $TypeText \in E_{dynamic}$  then
         $E_{dynamic} \leftarrow \langle TypeText, Wait \rangle$ ;
    else if  $TypeNumber \in E_{dynamic}$  then
         $E_{dynamic} \leftarrow \langle TypeNumber, Wait \rangle$ ;

    return  $E_{dynamic}$ ;

```

Algorithm 2 Dynamic Event Extraction


inferring specific event parameters automatically. For instance, if a DragSprite event is chosen due to a  block, the target sprite's position can precisely be determined from the program state. An overview of all events classified into inferable and non-inferable parameters is shown in Table 3. Since the dynamicEventExtraction extracts information about the current state of the program, the execution of the program has to be interleaved with the selection of events, as shown in Algorithm 3. Furthermore, because we assume the

Table 3 Inferable and non-inferable parameters mapped to their respective events

Event	Inferable parameter	Non-inferable parameter
KeyPress	Key	Duration
ClickSprite	Sprite	—
ClickStage	—	—
TypeText	Text	—
TypeNumber	—	Number
MouseDown	Press-status	—
MouseMove	—	(x , y) coordinates
MouseMoveTo	Sprite	—
DragSprite	Sprite/colour/edge	Angle
Sound	Volume, fixed duration	—
Wait	—	Duration

Greenflag event to be similar to the initiation of a regular program execution, the event is always sent at the beginning of test execution and never after the test execution has started.

```

input   : program  $p$  for which to generate tests
input   : number  $N \in \mathbb{N}$  of tests to generate
output  : test suite  $T$ 
function randomTestGenerationdynamic( $p, N$ ):
     $\gamma \leftarrow$  start  $p$  by sending Greenflag and obtain concrete program state;
     $T \leftarrow \langle \rangle$ ;
    for  $i \leftarrow 1$  to  $N$  do
         $E_{\text{dynamic}} \leftarrow$  dynamicEventExtraction( $\gamma$ );
         $e \leftarrow \text{random}(E_{\text{dynamic}})$ ;
         $\text{params} \leftarrow$  parameter list of  $e$ ;
         $\langle \text{params}_1, \dots, \text{params}_n \rangle \leftarrow$  inferable parameters in  $\text{params}$ ;
         $v \leftarrow \langle \rangle$ ;
        for  $j \leftarrow 1$  to  $n$  do
             $V_j \leftarrow$  value domain for  $\text{params}_j$ ;
             $v \leftarrow v \mathbin{++} \langle \text{random}(V_j) \rangle$ ;
         $T \leftarrow T \mathbin{++} \langle (e, v) \rangle$ ;
         $\gamma \leftarrow$  send  $(e, v)$  to  $p$ ;
    return  $T$ ;

```

Algorithm 3 Random test generator using dynamic event extraction

4.3 Assertion Generation

The test generation algorithm produces tests that consist of sequences of events. In order to be able to detect faults, a WHISKER test case also requires an *observer* (Section 2.2) that checks the observed behaviour against the expected behaviour. Since we envision that a common application scenario for WHISKER is that tests are generated on a model solution and then executed on student solutions, we operationalise the WHISKER observer in terms of *regression test assertions* that capture the state of the model solution. An assertion is a Boolean function that takes the program state P as input, and checks one of the properties against an expected value. If the value deviates, the assertion fails the test case. The assertions implemented in WHISKER are listed in Table 4: Each assertion is implemented in terms of code to check the value at runtime, and can synthesise JavaScript code that implements the observer in the generated WHISKER test.

The assertion generation algorithm is based on the approach proposed by Xie (2006), which essentially adds an assertion for every attribute after every step of a test, with expected values derived from the current version of the program. Since the number of possible assertions for a WHISKER test is proportional to the number of sprites, clones, and events in a test case, a direct application of the approach by Xie (2006) would lead to a huge number of assertions, many of which would be irrelevant for the program under test. While a common approach to filter relevant assertions is using mutation analysis (Fraser and Zeller 2011), the long test execution times of WHISKER tests together with the many executions required by the mutation analysis render this approach impracticable for WHISKER tests. We therefore reduce assertions as follows: We execute each test event by event; after each event, we determine the values for all possible assertions in the current state, and compare them against the values of the same assertions in the previous state. Only if the value of an assertion changes from the previous to the current state, this assertion is added to the current state.

Table 4 Types of assertions supported by the WHISKER test generator

Assertion	Target	Expected value
Backdrop	Stage	Costume (=backdrop) number
Clone count	Sprite	Number of clones
Costume	Sprite, Clone	Costume number
Direction	Sprite, Clone	Heading ± 1 degree
Graphics effect	Sprite, Clone, Effect type	Value
Layer	Sprite, Clone	Layer number
List	Sprite, Stage	Length
Position	Sprite, Clone	(x , y) coordinates ± 5 pixels
Say	Sprite, Clone	Existence, type, and contents of speech bubble
Size	Sprite, Clone	Size attribute
Variable	Sprite, Stage	Variable value
Visibility	Sprite, Clone	Visibility attribute
Volume	Sprite, Clone	Volume attribute

5 Test Generation Algorithms

5.1 Encoding SCRATCH Tests Using Grammatical Evolution

A prerequisite for applying search algorithms is a representation amenable to the modifications required by different search operators. For our application context we aim to evolve test cases, which consist of sequences of events. One challenge that applies to such sequences is that there can be dependencies between different events within the sequence. For example, assume two successive click events on the same sprite, where the execution of the first click event causes the sprite to be hidden—since the sprite is hidden, no second click can be performed on the sprite. Consequently, events cannot be performed in arbitrary order, and search operators that manipulate events may lead to invalid sequences. Rather than directly encoding test cases as sequences of events we therefore use an encoding inspired by Grammatical Evolution (O'Neill and Ryan 2001), where the mapping from genotype to phenotype is performed using a problem-specific grammar $G = \langle T, N, P, n_s \rangle$. Here, T is a set of terminals, which are the items that will appear in the resulting phenotype; N are non-terminals, which are intermediate elements associated with the production rules $P : N \rightarrow (N \cup T)^*$; the element $n_s \in N$ is the start symbol, which is used at the beginning of the mapping process.

The genotype is typically represented as a list of bits or integers; we use an integer representation (*codons*). Since a codon can represent the next event that will be executed or determine a non-inferable parameter of an event, we differentiate between *event codons* c_e and *parameter codons* c_p .

The mapping of a list of codons to the phenotype creates a derivation of the grammar as follows: Beginning with the first production of starting symbol n_s of the grammar, for each non-terminal x on the right-hand side of the production we choose the r th production rule for x . Given an event codon c_e and n productions for non-terminal x , the number r of the production rule to choose is determined as follows:

$$r = c_e \bmod n \quad (1)$$

In our application scenario, n represents the number of available events (Section 4). The resulting number r is then used as index for selecting an event from the available events E . If a look-up of Table 3 indicates that the selected event contains j non-inferable parameters, the subsequent j parameter codons c_p of the genotype are then queried to determine the required parameters based on a unique parameter mapping for each event type ev .

Note that a single change of one event codon might result in the selection of a different event that consumes more or less non-inferable parameters than the previous event. Such a change of consumed parameter codons would potentially lead to a diverging interpretation of the remaining codons because former event codons might be treated as parameter codons and vice versa. Thus, mutating a single event codon could result in an entirely different test execution, leading to considerable jumps within the fitness landscape. To avoid this problem, we assign each event codon a fixed number of parameter codons n_p , regardless of the number of non-inferable parameter codons a given event consumes. The fixed number of assigned parameter codons n_p is defined by the maximum number of consumed codons across all processable events of a given project E_{static} . A genotype is then incrementally translated into a phenotype by transforming each event codon to the corresponding SCRATCH event using (1) and consuming the event-specific number of required event parameter $j \leq n_p$.

Algorithm 4 describes the simultaneous decoding and execution of a codon sequence and demonstrates how concrete parameter values are chosen. The DragSprite event constitutes a special case since the dragging location defined by the event extractor tends to show unintended side effects. Consider, for example, a game in which the player has to reach a specific position at the end of a maze without touching a wall. If the DragSprite event moves the player sprite to the goal position in order to trigger code related to winning the game, the player sprite might also overlap with a wall, leading to a Game Over state instead. To avoid these side effects, the DragSprite event consumes an additional parameter codon to slightly move the determined dragging location in the direction of the parameter codon value, which is interpreted as an angle in the range of $[0, 360]$. Moreover, for KeyPress and Wait events, the user can specify an upper bound for the respective wait or keypress duration.

Each time an event has been inferred from the current event codon c_e , the respective event is sent to the SCRATCH VM, which updates its state based on the received event. Then, the decoding of the codon sequence moves on to the next unused event codon of the genotype, skipping all unused parameter codons c_p of the previous event codon. Overall, we define an implicit grammar where the starting production for a test case of length $l = \text{len}(\text{codons})/(n_p + 1)$ is given by the following formula:

$$\text{testcase} ::= \text{input}_1 \text{ input}_2 \dots \text{input}_l$$

Consider the following example chromosome that was generated for the program depicted in Fig. 8, with $E_{static} = \langle \text{Wait}, \text{ClickSprite}(\text{Cat}), \text{KeyPress}(\text{Space}) \rangle$ and $n_p = 1$:

$$T = \langle [4 \ 3] [5 \ 8] [2 \ 9] \rangle$$

For a better visualisation of the codon groups, we placed each event codon together with its n_p parameter codons inside rectangular brackets. After sending the Greenflag event to the SCRATCH VM, the resulting initial program state provides a choice of two events: a Wait event as well as a ClickSprite event on the cat sprite. Using the first event codon value 4 and the two available events $\langle \text{Wait}, \text{ClickSprite}(\text{Cat}) \rangle$, we compute $4 \bmod 2 = 0$ and thus select the Wait event from the set of available events. Since the Wait event requires a parameter that denotes the duration, the next codon 3 is interpreted as the event's parameter, i.e., as the number of steps to wait for. Moving on, the next event codon is 5. In our

```

input : program  $p$ 
input : list  $\langle c_1, \dots, c_l \rangle$  of codons
input : event extractor  $\text{extractEvents}$ 
input : user configuration  $u$ 
output : test case  $t$ 
function  $\text{decodeAndExecute}(p, \langle c_1, \dots, c_l \rangle, \text{extractEvents})$ :
     $t \leftarrow \langle \rangle$ ;  $i \leftarrow 0$ ;
     $\text{UpperBoundKeyPressDuration} \leftarrow u$ ;
     $\text{UpperBoundWaitDuration} \leftarrow u$ ;
     $n_p \leftarrow \text{maxConsumedCodons}(E_{\text{static}})$ ;
     $\gamma \leftarrow \text{start } p \text{ by sending Greenflag and obtain concrete program state}$ ;
    while  $i < l$  do
         $E \leftarrow \text{extractEvents}(\gamma)$ ;
         $n \leftarrow |E|$ ;
         $e \leftarrow \text{pick element } c_i \bmod n \text{ from } E$ ;
        if  $\text{eventType}(e) = \text{KeyPress}$  then // Determine key press duration
             $\text{duration} \leftarrow c_{i+1} \bmod \text{UpperBoundKeyPressDuration}$ ;
             $\text{setKeyPressDuration}(\text{duration})$ ;
        else if  $\text{eventType}(e) = \text{MouseMove}$  then // Determine x and y position
             $x \leftarrow (c_{i+1} \bmod \text{StageWidth}) - (\text{StageWidth}/2)$ ;
             $y \leftarrow (c_{i+2} \bmod \text{StageHeight}) - (\text{StageHeight}/2)$ ;
             $\text{setMouseMovePosition}(x, y)$ ;
        else if  $\text{eventType}(e) = \text{DragSprite}$  then
             $\text{angle} \leftarrow c_{i+1}$ ;
            if  $\text{angle} < 360$  then // Move x & y in the direction of angle
                 $\text{Sprite} \leftarrow \text{getSpriteToDrag}()$ ;
                 $(x, y) \leftarrow \text{getDragLocation}()$ ;
                 $x \leftarrow x + \text{horSize}(\text{Sprite}) * \cos(\text{angle})$ ;
                 $y \leftarrow y + \text{vertSize}(\text{Sprite}) * \sin(\text{angle})$ ;
                 $\text{setDragLocation}(x, y)$ ;
            else if  $\text{eventType}(e) = \text{Wait}$  then // Determine wait duration
                 $\text{duration} \leftarrow c_{i+1} \bmod \text{UpperBoundWaitDuration}$ ;
                 $\text{setWaitDuration}(\text{duration})$ ;
         $t \leftarrow t ++ \langle e \rangle$ ; // Add event including its parameters to test case
         $\gamma \leftarrow \text{send } e \text{ to } p$ ; // Send event including its parameters to SCRATCH VM
         $i \leftarrow i + n_p + 1$ 
    return  $t$ ;

```

Algorithm 4 Decoding and execution of a codon sequence

case, waiting does not have an impact on the list of available events, which means we have $\langle \text{Wait}, \text{ClickSprite}(\text{Cat}) \rangle$ as our sequence of available events again. Due to $5 \bmod 2 = 1$, we choose to click on the cat sprite. Since the ClickSprite event does not require additional parameter codons, we skip the reserved parameter 8. After executing the ClickSprite event, the event handler script broadcasts a message that triggers the receiver script in the bear sprite, which in turn waits for a press of the space key in an infinite loop. Thus, when interpreting the next event codon (2), there are three possible events to choose from: $\langle \text{Wait}, \text{ClickSprite}(\text{Cat}), \text{KeyPress}(\text{Space}) \rangle$. Because $2 \bmod 3 = 2$, the chosen event is to press the space key. Because the KeyPress event requires a single additional codon to determine the press duration, the following reserved parameter codon is consumed to define a press duration of 9. In order to explore the search space and derive new chromosomes, we apply mutation and crossover.

5.1.1 Mutation

During mutation, event codons are grouped with their n_p assigned parameter codons, which results in $n_g = \text{len}(\text{codons})/(n_p + 1)$ codon groups. Codon groups are then traversed and mutated with a probability of $1/n_g$. If a codon group is mutated, a single mutation operator out of the following three operators is chosen randomly. Each description of a mutation operator succeeds an exemplary mutant that results from applying the respective mutation operation at the codon group [5, 8] of the genotype $T = \langle [4\ 3]\ [5\ 8]\ [2\ 9] \rangle$.

- Add a novel codon group in front of the selected codon group, consisting of $n_p + 1$ randomly generated codon values: $T = \langle [4\ 3]\ [1\ 7]\ [5\ 8]\ [2\ 9] \rangle$
- Iteratively modify every codon value of the selected codon group by sampling new codon values from a gaussian distribution that has its mean value set to the respective codon value: $T = \langle [4\ 3]\ [4\ 10]\ [2\ 9] \rangle$
- Delete the selected codon group: $T = \langle [4\ 3]\ [2\ 9] \rangle$

Since each codon group is mutated with a probability of $1/n_g$, we apply, on average, one mutation operation to every parent.

5.1.2 Crossover

Crossover takes as input two parents and produces two children by combining the codons of both parents at specific codon positions ψ_1 and ψ_2 . Similar to the mutation operator, the crossover operator starts by grouping the codons of both parents into codon groups of sizes n_{g1} and n_{g2} . To derive the crossover positions ψ_1 and ψ_2 for both parents, we first randomly select a relative crossover point ψ_r in the range of $[0, 1]$, with 0 representing the first codon group and 1 representing the last codon group of any given parent, and map ψ_r to the corresponding codon group for each parent individually. A new offspring is generated by combining the first parent's codon groups from 0 to $\psi_1 - 1$ with the codon groups residing at the index positions ψ_2 to n_{g2} of the second parent. Finally, a second child is produced by swapping the first and second parent. For example, consider the two parents

$$\begin{aligned} T_1 &= \langle [0\ 1]\ [2\ 3]\ [4\ 5]\ [6\ 7]\ [8\ 9] \rangle \\ T_2 &= \langle [10\ 11]\ [12\ 13]\ [14\ 15] \rangle \end{aligned}$$

and a randomly chosen relative crossover point $\psi_r = 0.5$, which is mapped to the crossover positions $\psi_1 = 2$ and $\psi_2 = 1$. Then, using the derived crossover positions, the crossover operator produces the two children:

$$\begin{aligned} T_{12} &= \langle [0\ 1]\ [2\ 3]\ [12\ 13]\ [14\ 15] \rangle \\ T_{21} &= \langle [10\ 11]\ [4\ 5]\ [6\ 7]\ [8\ 9] \rangle \end{aligned}$$

5.2 Fitness Function

Fitness functions offer a means of distinguishing “good” and “bad” chromosomes. We aim for statement coverage, such that the fitness function estimates how close a given execution was to reaching a statement. This estimate is traditionally calculated by considering distances in the control flow (approach level (Wegener et al. 2001)) and distances of the executions of individual conditional statements (branch distance (Korel 1990)):

- Approach Level: A target statement can be nested arbitrarily deep in conditional (e.g., if-then-else) parts of the program. The control flow only reaches the target if it takes

specific branches at these decision nodes. Intuitively, the approach level measures the minimum number of decision nodes by which control flow missed the target.

- **Branch Distance:** If the control flow takes the wrong branch at any of the dependent decision nodes, then the branch distance estimates how close the underlying conditional statement was from taking the opposite branch.

The fitness calculation in WHISKER is based on these concepts, but requires adaptations.

5.2.1 Interprocedural Control Flow and Control Dependence Graphs

The approach level metric was designed for procedural code containing potentially deeply nested code constructs. In contrast, SCRATCH programs tend to consist of many small scripts, which communicate through events and messages. To counter this discrepancy, we create an interprocedural control flow graph and control dependency graph for SCRATCH programs, and use this for calculating approach levels.

A given target SCRATCH program consists of a number of scripts; for each script we derive the control flow graph (CFG), defined as $CFG = (L \cup \{entry, exit\}, G)$, i.e., a directed graph consisting of control flow locations L as well as dedicated *entry* and *exit* nodes, and edges based on the control flow G between these nodes. We combine these intraprocedural CFGs to an interprocedural super-CFG as follows:












- For each event handler , we add an artificial node with edges to the event handler (*hat block*) as well as the *exit* node. We further add an edge from *entry* to this artificial node for event handlers of user inputs.
- For each  or  statement, we add an edge from the broadcast to all scripts that start with a matching receive event handler block .
- For each  statement, we add an edge to all scripts that start with a matching  event handler block for the corresponding sprite.
- For each  and  statement, we add an edge from that block to all scripts that start with a matching  block for the corresponding backdrop (or all such event handlers if the name of the backdrop is not known).
- For each *procedure call* statement , we add an edge from the call to the start block of the procedure (custom block), and a return edge from its end to the successor node in the calling script. If there are multiple calls to a custom block, all calls lead to the same start block, and there are multiple return edges from the end of the procedure.
- For  we add an edge to the exit node, since the rest of the execution is dependent on the condition being satisfied.

Figure 9a shows the interprocedural CFG for the program in Fig. 8. This CFG contains two artificial event nodes (*clicked cat?*, *broadcast?*), each of which effectively is a branching statement depending on whether the event occurs. These branches turn the occurrence of events into control dependencies of the statements in the event handler code.

5.2.2 Approach Level

The approach level is calculated using the control dependence graph, which can be derived directly from the interprocedural CFG. Figure 9b shows the CDG for the example program

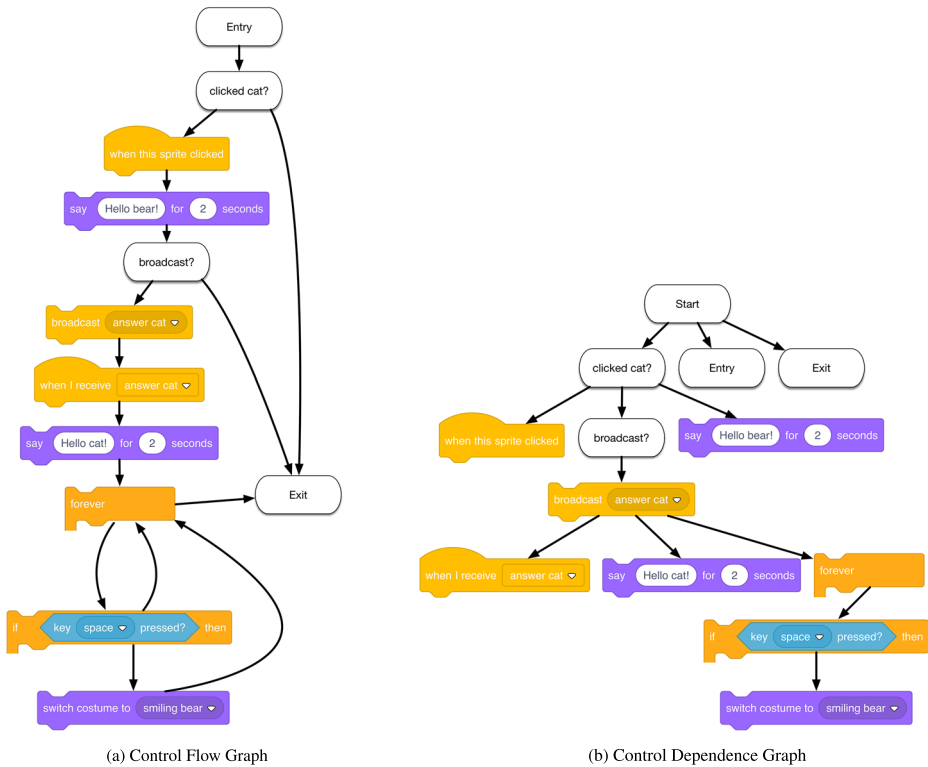


Fig. 9 Interprocedural control flow graph and control dependence graph created for the example program from Fig. 8

shown in Fig. 8; inter-dependencies, for example caused by message broadcasts, are captured in this graph. Figure 10 shows a simple example program, where the say block is only executed if the two control dependencies checking x against 50 and then against 60 are satisfied. If x is less than 50, then the approach level after executing this script will be 1; if x is larger than 50 but less than or equal to 60, then the approach level is 0.

To measure the approach level, each statement fitness function pre-computes the distance in the CDG for each node to the target node. We instrument program executions by extending the SCRATCH VM such that information about executed branching statements is added to execution traces. Execution traces consist of block traces, which collect specific block-related data such as the block type and argument values of conditional statements. Given an execution trace, we iterate over the covered nodes and determine the minimum distance observed along the trace using the pre-computed approach levels.

5.2.3 Control Flow Distance

Traditional programs execute their procedures from start to end. In contrast, SCRATCH programs tend to execute for long durations because of the animations they tend to contain, while WHISKER tests at the same time impose strict time limits (defined implicitly by the number and duration of the Wait events in a test). When a WHISKER test reaches its end it is terminated, which may interrupt the execution at any point in the control flow. Such an




Fig. 10 Example program to illustrate aspects of the fitness function

interrupted execution might be following the correct path in the control flow, such that the branch distance of those executed control dependencies is 0 (i.e., the correct branch was taken). The search would now receive no guidance towards reaching the next control dependency or the target statement. In order to counter this problem, we introduce the control flow distance metric, which informs the fitness function how close an execution was to reaching a target node within a sequence of statements; the target node might either be the target of the fitness function itself, or the next control dependency on the path to it.

The computation of the control flow distance is outlined in Algorithm 5: Given a CFG G with control locations L , and a set of already covered locations $C \subset L$, we perform a *backwards* breadth-first search starting from the desired target node $t \in L$, and compute the minimal distance between t and any covered statement.

The function is called using the coverage information contained in a trace. The target node passed as parameter is the target node itself if the approach level is 0. Otherwise, we determine the control flow distance towards each successor control dependency, and select the minimum value. In Fig. 10, if the variable x is larger than 60, the execution may still be interrupted before the *say* block is executed. Therefore, once the second if-condition has been evaluated, the control flow distance to the *say*-block is 2, and after the *change* statement it is 1.

5.2.4 Branch Distance Instrumentation

The branch distance estimates how close a conditional statement was to evaluating to a specific outcome (true or false). We extended the SCRATCH VM such that for each conditional statement the branch distances are calculated and traced. For each conditional statement, the execution trace contains information about the minimum branch distances (for evaluation to true and to false) observed during an execution. To calculate the branch distance we first select the closest control dependence, as determined when calculating the approach

```


input   : A CFG  $G$  with control locations  $L$ 
input   : Set of covered locations  $C \subset L$ 
input   : Target node  $t \in L$  for which to compute the control flow distance
output  : The minimal distance  $d \in \mathbb{N}$  between  $n$  and a covered node in  $C$ 
function controlFlowDistance( $G, C, t$ ):
     $d \leftarrow 0$ ;                                     // initialise distance
     $Q \leftarrow \langle (t, d) \rangle$ ;                       // initialise work queue
     $V \leftarrow \{t\}$ ;                                // initialise set of visited nodes
    while  $Q \neq \langle \rangle$  do
         $\langle (n, d), \dots \rangle \leftarrow Q$ ;                // get first element of  $Q$ 
         $Q \leftarrow Q \setminus \langle (n, d) \rangle$ ;
        if  $n \in C$  then
            return  $d$ ;
         $V \leftarrow V \cup \{n\}$ ;
         $P \leftarrow$  immediate predecessors of  $n$  in  $G$ ;
        foreach predecessor  $p \in P$  do
            if  $p \notin V$  then
                 $Q \leftarrow Q \cup \langle (p, d+1) \rangle$ ;
    return  $d$ ;


```

Algorithm 5 Control flow distance computation via breadth-first traversal



level, and then select the minimum branch distance of the outgoing edge that would take the execution closer to the target node.

Suppose the target is to reach the *say* block in Fig. 10, then if x is, for example, 42, then the branch distance is computed based on the first if-condition as $|50 - 42| + 1 = 9$. If the first if-condition evaluates to true, for example with $x = 55$, then the second if-condition serves to calculate the branch distance $|60 - 55| + 1 = 6$.




The instrumentation applies the regular equations known from the literature (Korel 1990); for example, given an equality comparison , the distance for this condition to evaluate to true is 0 if x equals 42 and otherwise $|x - 42|$; the distance for this condition to evaluate to false is 0 if x is already different from 42, otherwise it is 1. The instrumentation of the SCRATCH VM implements this for all standard relational and logical operators.

Due to the game-like nature of many SCRATCH programs, a common task is to check for interactions between sprites, e.g. whether they are touching. To this end, SCRATCH provides dedicated *sensing blocks*. These can be used as conditions for if-then-else or loop blocks, but are also often found in combination with  blocks, which are encoded as branching statements in the CFG. Therefore, branch distance needs to be calculated and traced for all sensing blocks, but the equations presented for standard operators cannot be applied, and we require a novel definition of branch distance for sensing blocks.

A dichotomous notion (e.g., using 0 or 1 depending on whether a sensing block reports true or false) leads to challenging plateaus in the fitness landscape, which reduces the effectiveness of the search. This is a well-known problem to the test generation community. It has been shown that altering the fitness landscape and restoring lost gradients can lead to better guidance and consequently improvements in the search (Vogl et al. 2021).

The key to this is the observation that many sensing blocks query the location of objects on the stage. For example,  or  blocks check if the current sprite or one of its colours is touching the other target colour. We transform these conditions by checking if the *Euclidean distance* between the subjects is 0 on the canvas, and use it as the






branch distance. This fits the traditional notion: if the condition is true, both the Euclidean distance and the branch distance to the true-branch are 0, and the distance to the false-branch is 1. On the other hand, if the condition is false, we define the distance to the false-branch as 0, and use the Euclidean distance for the true-branch. This way, sprites or colours that are closer together will have smaller distance values than the ones that are further apart.

Similarly, if the condition checks if a sprite is touching the edge of the stage () , we can gather the position information and calculate the distance to all four edges, and use the minimal distance as the branching distance. If the condition checks if a sprite is  or  , we use the distance between the sprite and the mouse pointer or the target sprite, respectively, as the branch distance.

The repeat-times block also represents a special case since it does not evaluate a condition expressed in code. We therefore instrumented these loops such that the branch distance to exiting the loop is represented by the remaining number of loop iterations, and the false distance is only non-0 when the loop is exited. To ensure that these loops are considered as control dependencies, we add an edge in the CFG from the loop header to the exit node.

The occurrence of events (e.g., greenflag, key press, sprite click) is encoded in artificial branching nodes in the CFG (cf. Section 5.2.1). The occurrence or absence of events is encoded with branch distances of 1 or 0, depending on whether or not the event occurred.

5.2.5 Time-Dependent Statements

SCRATCH contains several time-dependent statements, such as explicit waits  , think/say blocks  , glide-animations  , audio playback with a certain duration  , or text to speech translation  . The control flow distance only provides limited guidance in terms of the number of statements remaining to be covered in a sequence. To better capture time in the fitness function, we model time-dependencies explicitly in the CFG, and include information on remaining times in the branch distance.

To achieve this, we add artificial edges for each time-dependent statement to the exit node in the CFG, turning them into control dependencies of all successor statements, thus including them in the calculation of the approach level. We extend the instrumentation of the SCRATCH VM such that traces include branch distances for time-dependent statements: If a time-dependent statement was fully executed, then the true distance is 0 and the false distance is 1; if the execution was interrupted before the statement completed, then the true distance is defined as the remaining time (and the false distance is 0).

5.2.6 Overall Fitness Function

The overall fitness function for a specific target node is a combination of Approach Level, Branch Distance and Control Flow Distance: The approach level value is an integer, and to avoid creating deceiving fitness landscapes it needs to dominate the other measurements. We therefore normalise branch distance and control flow distance in the range [0, 1] using the normalisation function $\alpha(x) = x/(1+x)$ (Arcuri 2013). If the branch distance is greater than 0, then we set the control flow distance to the maximum value of 1, since the execution first needs to change the evaluation of the last control dependency before progressing in the CFG matters. To ensure the dominance of the approach level, we thus need to multiply it by a factor of 2. Finally, we determine the fitness $f = \text{fitness}(t) \in \mathbb{R}$ of a test t for a given

target location as described in Algorithm 6. Since our overall objective is to achieve full coverage, one such fitness function is created for each block in the program.

```

input   : A test case  $t$ 
input   : Control flow graph  $G$  (implicitly)
input   : Target location  $s$  (implicitly)
output  : The fitness for  $t$ 
function fitness( $t$ ):
     $C \leftarrow$  set of locations in  $G$  covered by  $t$ ;
     $a \leftarrow 2 \times \text{approachLevel}(G, C, s)$ ;
     $b \leftarrow \alpha(\text{branchDistance}(G, C, s))$ ;
    if  $b > 0$  then
         $c \leftarrow 1$ ;
    else
         $c \leftarrow \alpha(\text{controlFlowDistance}(G, C, s))$ ;
    return  $a + b + c$ ;

```

Algorithm 6 Fitness computation for test cases

5.3 Search Algorithms

Given the encoding and fitness function, it is possible to apply any meta-heuristic search algorithm to the problem of test generation for SCRATCH. Random search (Algorithm 7) is

```

input   : SCRATCH program  $p$ 
output  : Test suite  $T$ 
function randomSearch( $p$ ):
     $G \leftarrow$  coverage goals in  $p$ ;                                // Set of all goals
     $C \leftarrow \emptyset$ ;                                           // Set of already covered goals
     $T \leftarrow \emptyset$ ;                                           // Test suite

    while stopping condition not reached do
         $\langle c_1, \dots, c_k \rangle \leftarrow \text{generateRandomCodons}(p)$ ;
         $t \leftarrow \text{decodeAndExecute}(p, \langle c_1, \dots, c_k \rangle, \text{extractEvents})$ ;
         $C' \leftarrow$  goals covered by  $t$ ;
        if  $C' \setminus C \neq \emptyset$  then                                // Check if  $t$  covers new goals
             $T \leftarrow T \cup \{t\}$ ;
             $C \leftarrow C \cup C'$ ;

    return  $T$ ;

function generateRandomCodons( $p$ ):
     $\bar{C} \leftarrow \langle \rangle$ ;
     $i \leftarrow 0$ ;
     $k \leftarrow \text{random}(\{1, 2, \dots\})$ ;
    while  $i < k$  do
         $c \leftarrow \text{random}(\{0, 1, \dots\})$ ;
         $\bar{C} \leftarrow \bar{C} \cup \{c\}$ ;
         $i \leftarrow i + 1$ ;
    return  $\bar{C}$ ;

```

Algorithm 7 Random search

the simplest conceivable global search algorithm. As a global search algorithm, it considers the *entire* search space (in contrast to neighbouring chromosomes in a local search algorithm), trying to cover as many statements at a time as possible. It operates by repeatedly sampling a test t at random, and adds it to the test suite T if it covers a new target. This process continues until a given search budget is exhausted, after which T is returned. Due to its simplicity, random search is often used as a baseline for comparison. As many SCRATCH programs tend to be small, it is also possible that random search will often be sufficient in order to generate adequate test suites.

The aim of automated test generation is to maximise the achieved code coverage, which is a task that lends itself to a multi-objective problem representation, where every single statement is an individual optimisation goal in its own right. Hereby, it is not uncommon to encounter conflicting goals (e.g., statements in `if`-branches vs. `else`-branches), and depending on the size of p , the number of goals might range from tens to hundreds, or possibly even more. This poses scalability challenges to “traditional” many-objective algorithms,

```

input   : program  $p$  for which to generate tests
input   : population size  $N \in \mathbb{N}$ 
output  : test suite  $A$  (an “archive” of test cases)
function MOSA( $p, N$ ):
     $P \leftarrow \text{randomPopulation}(p, N)$ ;
     $T \leftarrow \text{decodeAndExecuteAll}(p, P)$ ;
     $A \leftarrow \text{updateArchive}(\emptyset, T)$ ;
    while stopping condition not satisfied do
         $O \leftarrow \text{generateOffspring}(P)$ ;
         $T \leftarrow \text{decodeAndExecuteAll}(p, O)$ ;
         $A \leftarrow \text{updateArchive}(A, T)$ ;
         $\langle F_0, F_1, F_2, \dots \rangle \leftarrow \text{preferenceSorting}(P \cup O)$ ;
         $P \leftarrow \emptyset$ ;
         $i \leftarrow 0$ ;
        while  $|P| + |F_i| \leq N$  do
             $P \leftarrow P \cup F_i$ ;
             $i \leftarrow i + 1$ ;
         $\langle c_1, c_2, \dots \rangle \leftarrow \text{sortByCrowdingDistance}(F_i)$ ;
         $P \leftarrow P \cup \{c_j \mid j = 1, 2, \dots, (N - |P|)\}$ ;
         $(P, A) \leftarrow \text{localSearch}(P, A)$ ;
    return  $A$ ;

function randomPopulation( $p, N$ ):
     $P \leftarrow \emptyset$ ;
     $i \leftarrow 0$ ;
    while  $i < N$  do
         $c \leftarrow \text{generateRandomCodons}(p)$ ;
         $P \leftarrow P \cup \{c\}$ ;
    return  $P$ ;

function decodeAndExecuteAll( $p, C$ ):
     $T \leftarrow \emptyset$ ;
    foreach codonSequence  $c \in C$  do
         $t \leftarrow \text{decodeAndExecute}(p, c, \text{extractEvents})$ ;
         $T \leftarrow T \cup \{t\}$ ;
    return  $T$ ;

```

Algorithm 8 Many-Objective Sorting Algorithm (Panichella et al. 2015)

```

input : program  $p$ 
input : set  $K$  of targets in  $p$ 
input : start of the focused phase  $F \in [0, 1]$ 
input : initial population size per target  $n_0 \in \mathbb{N}$ 
input : final population size per target  $n_f \in \mathbb{N}$ 
input : initial random selection probability  $r_0 \in [0, 1]$ 
input : final random selection probability  $r_f \in [0, 1]$ 
input : initial maximum mutation count  $m_0 \in \mathbb{N}$ 
input : final maximum mutation count  $m_f \in \mathbb{N}$ 
output : test suite  $T$ 
function MIO( $p, K, F, n_0, n_f, r_0, r_f, m_0, m_f$ ):
     $\bar{A} \leftarrow \emptyset^{|K|}$ ;
     $B \leftarrow$  depleted search budget;
     $n \leftarrow n_0$ ;
     $r \leftarrow r_0$ ;
     $m \leftarrow m_0$ ;
    while stopping condition not satisfied do
        if  $B < F$  then
             $n \leftarrow \text{calculateDynamicParameter}(n_0, n_f, B, F)$ ;
             $r \leftarrow \text{calculateDynamicParameter}(r_0, r_f, B, F)$ ;
             $m \leftarrow \text{calculateDynamicParameter}(m_0, m_f, B, F)$ ;
        if  $\bar{A} = \emptyset^{|K|}$  or with probability  $r$  then
             $\langle c_1, \dots, c_k \rangle \leftarrow \text{generateRandomCodons}(p)$ ;
             $\text{decodeAndExecute}(p, \langle c_1, \dots, c_k \rangle, \text{extractEvents})$ ;
             $\bar{A} \leftarrow \text{updateArchive}(\bar{A}, t, n)$ ;
             $\bar{A} \leftarrow \text{localSearch}(\{t\}, \bar{A})$ ;
        else
             $k \leftarrow \text{chooseTarget}(K)$ ;
             $t \leftarrow \text{random}(A_k)$ ;
             $i \leftarrow 0$ ;
            while  $i < m$  do
                 $t' \leftarrow \text{mutate}(t)$ ;
                 $\bar{A} \leftarrow \text{updateArchive}(\bar{A}, t', n)$ ;
                if  $t'$  has a better fitness value than  $t$  for target  $k$  then
                     $t \leftarrow t'$ ;
                 $i \leftarrow i + 1$ ;
             $\bar{A} \leftarrow \text{localSearch}(\{t\}, \bar{A})$ ;
         $B \leftarrow$  updated search budget;
     $T \leftarrow \text{constructTestSuite}(\bar{A}, K)$ ;
    return  $T$ ;

function calculateDynamicParameter( $x_0, x_f, B, F$ ):
    return  $x_0 + (x_f - x_0) \cdot \frac{B}{F}$ ;

```

Algorithm 9 Many Independent Objective Algorithm (Arcuri 2017)

such as the well-known NSGA-II (Deb et al. 2002). Due to the so-called dominance resistance phenomenon, the proportion of non-dominated solutions increases *exponentially* with the number of goals to optimise, thus degrading the search to a random one in the process.

We therefore employ the Many-Objective Sorting Algorithm (MOSA) (Panichella et al. 2015) (outlined in Algorithm 8). MOSA is a modified variant of NSGA-II that caters to mentioned peculiarities of test generation. Most notably, it introduces a so called *preference criterion* that allows us to assign a preference among non-dominated solutions based on how “close” they come to covering a *new, previously uncovered* target. This way, the number

of targets to consider at a time is reduced and the search budget is directed towards the targets that are still left to be covered. As a notable deviation from the original algorithm (Panichella et al. 2015), we have extended the search algorithm to a memetic algorithm (Fraser et al. 2015) by applying a local search to the new parent population at the end of every generation (and also updating the archive if necessary). A local search algorithm explores the local neighbourhood of a candidate solution in a more focused way than a global exploration would. In particular, this post-processing step is necessary to address the challenge of finding suitable test execution durations specific to SCRATCH.

The number of targets to cover in a SCRATCH program can be very large, and even though MOSA tries to address this problem, it might still struggle to achieve 100% coverage. For example, certain statements may be infeasible, and it is thus not worthwhile trying to cover them. For this reason, the Many Independent Objective algorithm (MIO) (Arcuri 2017), outlined in Algorithm 9, tries to strike a balance between exploration and exploitation by focusing on those goals that are most promising, given the resources available. It has been specifically designed for test generation and is based on the (1+1) EA evolutionary algorithm, but also maintains an archive of candidate solutions for each coverage objective. Similar to MOSA, we extended MIO with local search by applying the local search defined in Section 5.3 with a certain probability after each generation.

We have turned MOSA and MIO into memetic algorithms by integrating local search operators (Fraser et al. 2015), which take an existing test case as input and explore its *neighbours* by applying operator-specific changes to its genotype. If a local search operator manages to generate an improved test case, which is measured based on the pursued goal of the applied operator, the original test case is replaced with the modified one.

5.3.1 Extension Local Search

SCRATCH projects often contain blocks that pause the execution of a program for a certain amount of time (Section 2.1). Due to these, some program statements can only be reached after waiting for an extended period of time. Extension local search aims to overcome these execution halting blocks by repeatedly adding `WaitEvents` to a given test case, eventually making previously hard to reach blocks more accessible. In order to append and execute `WaitEvents` at the end of a test case, the operator first has to obtain the SCRATCH VM's state after executing the original test case. For this purpose, the extension local search algorithm, shown in Algorithm 10, starts with the re-execution of the original test case.

In the while-loop that follows, the operator repeatedly checks for the presence of `TypeTextEvents`/`TypeNumberEvents` and novel events that were not present in the list of events E during the previous iteration of the loop. If a `TypeTextEvent` or `TypeNumberEvent` is found, it is preferably selected over `WaitEvents` since the program execution halts until the user has given an answer to a posed question, which is indicated to the user through an UI-focus switch that highlights a text field. Hence, until an answer has been given in the form of a `TypeTextEvent`/`TypeNumberEvent`, adding additional `WaitEvents` usually does not contribute in exploring novel program states. Newly found events, on the other hand, are preferred because these events become available due to overcoming certain execution halting blocks, and therefore promise to lead to novel program states. However, newly discovered events are only selected with a specific probability, as otherwise, program states potentially hiding behind even longer wait durations would most likely remain out of reach. A prominent example of such a novel event scenario is a `ClickSpriteEvent`, representing a click on a button that gets only clickable after an introductory animation has finished.


```

input : list of codons  $\langle c_1, \dots, c_l \rangle$ 
input : event extractor extractEvents
input : program  $p$ 
input : user configuration  $u$ 
output : modified test case  $t'$ 
function extensionLocalSearch( $\langle c_1, \dots, c_l \rangle, p, \text{extractEvents}, u$ ):
    (MaxCodonLength, NewEventProbability)  $\leftarrow u$ ;
     $n_p \leftarrow \text{getNumParameterCodons}(t')$ ;
    done  $\leftarrow$  False;
     $j \leftarrow 0$ ;
     $t' \leftarrow \text{decodeAndExecute}(p, \langle c_1, \dots, c_l \rangle, \text{extractEvents})$ ;
     $\gamma \leftarrow \text{fetch program state of } p \text{ after executing } t$ ;
     $E_{prev} \leftarrow \text{extractEvents}(\gamma)$ ;
    while isActive( $\gamma$ ) and length( $c$ )  $\leq$  MaxCodonLength and  $\neg$ done do
         $E \leftarrow \text{extractEvents}(\gamma)$ ;
        // Apply TypeTextEvent/TypeNumberEvent.
        if TypeTextEvent  $\in E$  then
             $e \leftarrow \text{getRandomTypeTextOrNumberEvent}(E)$ ;
            // Apply a newly found event.
        else if  $E_{prev} \neq E$  and random( $\{0, \dots, 1\}$ )  $\leq$  NewEventProbability then
             $e \leftarrow \text{getRandomNovelEvent}(E)$ ;
            // Apply WaitEvent.
        else
             $e \leftarrow \text{WaitEvent}$ ;
         $c \leftarrow c \text{ ++ indexOf}(e, E) \text{ ++ generateParameterCodons}(n_p)$ ;
         $t' \leftarrow t' \text{ ++ } \langle e \rangle$ ;
         $\gamma \leftarrow \text{send } e \text{ to } p$ ;
        if fitnessHasNotImproved( $t'$ ) then
            done  $\leftarrow$  True;
         $E_{prev} \leftarrow E$ ;
    return  $t'$ ;

function generateParameterCodons( $n_p$ ):
     $i \leftarrow 0$ ;
     $c_p \leftarrow \langle \rangle$ ;
    while  $i < n_p$  do
         $c_p \leftarrow c_p \text{ ++ random}(\{1, 2, \dots\})$ ;
    return  $c_p$ ;

```

Algorithm 10 Extension Local Search

Furthermore, to save time the operator checks after every iteration if the overall fitness of the test case has improved and stops if no improvement could be observed. In addition to the lack of observable fitness improvements, the algorithm also stops if the maximum codon length defined by the user has been reached or if the program has stopped due to the execution of a  `stop` block or having reached the end of all program scripts s .

Since the extension local search operator aims to discover novel program states by extending the genotype of a test case, the operator can only be applied if a genotype has not reached its full length yet. Finally, the original test case is replaced with the extended one if the operator discovered new statements not previously covered by the existing test case.

5.3.2 Reduction Local Search

Reduction local search aims to reduce the codon length by removing genes that did not contribute to improving the fitness. For that purpose, every time a test case is executed, we save the codon position C_l that points to the last codon group after which no further fitness improvements have been observed. Using the index C_l , the reduction local search operator generates a new test case by cloning the codon groups located at the positions $[0, C_l]$, excluding all codons occurring after C_l . A significant difference to the extension local search operator is that reduction local search does not re-execute the original test case since C_l is already saved during the execution of the original test case.

The operator is only applied to genotypes for which $C_l < n_g$ is satisfied, with n_g representing the total number of codon groups contained in the genotype. Furthermore, since C_l points to the codon group after which no further fitness improvements have been observed, it is ensured that no covered statements are lost in the process of reducing the size of the genotype. Hence, every reduced test case is guaranteed to be an improvement over the original test case in terms of codon size and will therefore replace it in the search algorithm's population.

The benefit of applying reduction local search to a given test case is twofold: First, removing codon groups can save valuable search time due to not re-executing events that do not contribute in discovering new program states. Second, since the presented mutation operator mutates each codon group with a probability of $1/n_g$, reduction local search forces the mutation process to focus on relevant codon groups by increasing their mutation probability.

5.4 Test Minimization

Although MIO and MOSA both use minimization as a secondary criterion, the final test suite may contain test cases that are not minimal. We therefore apply a post-processing step that removes redundant events from test cases. The minimization algorithm produces test cases that are 1-minimal: A test case $T = \langle e_1, e_2, \dots, e_n \rangle$ of length n , where e_i can be interpreted as either a codon group or as an event, is 1-minimal with respect to a coverage goal represented by fitness function f , if for all i the test case $T' = \langle e_1, e_{i-1}, e_{i+1}, \dots, e_n \rangle$ has $f(T') > f(T)$. That is, removing any of the events leads to the test case no longer satisfying the coverage goal. We use the minimization algorithm implemented in EVOSUITE (Fraser and Arcuri 2012): For each test case T we iterate over all e_i starting from the last event, produce a test case T' without that event, and measure its fitness; if the fitness is not worse, then e_i is discarded and $T = T'$. In theory, a more efficient algorithm such as delta debugging could be used to increase the performance of the minimization (Leitner et al. 2007).

6 Experiments

To provide a better understanding of the problem of test generation for SCRATCH, we aim to answer the following questions:

- RQ1:** How much can test execution be accelerated reliably?
- RQ2:** Can SCRATCH projects be trivially covered?
- RQ3:** What is the best test generation algorithm for SCRATCH programs?
- RQ4:** How effective are generated tests at detecting faults?

6.1 Experimental Setup

6.1.1 Dataset 1: Projects with Manually Written WHISKER Tests (MANUAL)

The first dataset consists of SCRATCH programs with corresponding handwritten WHISKER tests from two previous studies. The first study (Greifenstein et al. 2021) provides 14 SCRATCH programs together with matching tests; the second study considered is the original WHISKER study (Stahlbauer et al. 2019), which provides the fruit catching game (Fig. 1) and its 28 test cases. Overall, the dataset comprises 15 programs with 94 manually crafted WHISKER test cases (11.3 on average per project) that contain a total of 185 assertions.

6.1.2 Dataset 2: Buggy and Correct Projects (BUGS)

Our second dataset consists of programs with bugs, and fixed versions thereof. The source is again the study by Greifenstein et al. (2021), in which students received the 14 programs included in our first dataset, but each with one intentionally seeded bug. The students then attempted to fix the bug, and the correctness of the student submissions was determined using the manually written WHISKER tests as well as manual evaluation by a teacher, who classified the programs into “correct” and “buggy” submissions. We create the dataset by collecting all student solutions that were evaluated by a human examiner. This process results in 559 manually reviewed SCRATCH programs (39.93 on average per project), of which 338 were rated as “correct” and 221 were rated as “buggy”.

6.1.3 Dataset 3: 1000 Random Projects (RANDOM1000)

SCRATCH is among the most popular platforms for programming beginners, and backed by a large online ecosystem and community. We created a dataset of publicly shared SCRATCH projects by mining projects as follows: Each SCRATCH project has a unique ID which is sequentially increasing as new projects are created. By probing project IDs we determined that starting roughly from ID 400.000.000 we can reliably retrieve projects in the format of SCRATCH 3, whereas below we frequently encountered version 2 projects. While WHISKER can also handle projects saved in version 2 of SCRATCH, our static analysis tool LITTERBOX (Fraser et al. 2021) used in related research on the same dataset (Adler et al. 2021) requires version 3 projects. We then uniformly sampled project IDs in the range of 400.000.000 and 700.000.000 (i.e., a number larger than the latest projects at the time of this writing), and downloaded batches of 1000 starting from the random ID using the REST-API provided by the SCRATCH webserver. Projects can only be downloaded if they are publicly shared. This process resulted in a dataset of 2.2 million projects, of which 1.500.937 are not remixes of other projects, i.e., variations of already uploaded projects. From these, we randomly sampled a subset of 1000, which is a compromise between a large desirable set of evaluation subjects and the resulting computational costs of running experiments in multiple different configurations with many repetitions to counter randomness.

Figure 11a shows the log-scale distribution of sizes of the 1000 projects based on the count of statement-blocks. Note that we only count executable statement blocks; i.e., we exclude any loose blocks, or blocks contained in dead code (e.g., event handlers for events not generated in the project). The majority of projects are small and have less than 100 statements, but there are also larger projects with up to 1082 (connected, reachable) statements.




Fig. 11 Distribution of number of statements per project

6.1.4 Dataset 4: Top 1000 Most Loved Projects (TOP1000)

Many of the perils of mining GitHub open source projects do not apply to our sampling process of SCRATCH projects: For example, whereas when mining GitHub projects it is often problematic that many projects are personal (i.e., there is only one committer, and collaboration with others is not intended) or do not even contain code (Kalliamvakou et al. 2016), all SCRATCH projects per definition are personal and contain SCRATCH code since that is the only use case. However, since a common application scenario for SCRATCH is creative use rather than programming it is possible that a random sample may contain many trivial projects. Thus, while the random sample is useful for external validity, we also created a second dataset of *popular* projects. Indeed it has been shown that a focus on the star-rating in GitHub projects can have an impact on resulting findings (Maj et al. 2021). To derive a dataset of popular projects, we crawled the <https://scratchstats.com/> website, which collects live statistics on SCRATCH users and projects, on 2021–12–16 and identified and downloaded the 1000 most “loved” projects on SCRATCH. Figure 11b shows that these projects are substantially larger, with sizes ranging from 5 to 11 530 connected and reachable statements, with a mean of 1036 statements. On average, these projects have received 10 315.6 loves and 432 589 views by other SCRATCH users. The dataset includes 21 projects which are a remix of other projects in the dataset.

6.1.5 Implementation and Tuning

The ideas presented in this paper are implemented as an extension to the test generation tool WHISKER (Stahlbauer et al. 2019). In particular, we added the SCRATCH VM modifications presented in Section 3, the event selection mechanisms outlined in Section 4 and the algorithms described in Section 5. The source code used in this study is publicly available ².

Every described test generation algorithm is accompanied by a set of configurable parameters that guide the search for a test suite. To optimise these parameters, we establish an optimisation dataset, which is disjoint from the TOP1000 and RANDOM1000 datasets, by randomly sampling 250 projects following the same procedure as in Section 6.1.3. Parameters used across all algorithms, and MOSA-specific ones, are determined by executing

²<https://github.com/se2p/whisker> based on the commit starting with 8537271, and <https://github.com/se2p/scratch-vm> based on the commit starting with ed4055f2

MOSA on all 250 projects using different values for a single parameter while fixating all other configurable parameters. After the search has finished, we compare the achieved coverages and choose the best performing configuration. Finally, we repeat the same procedure with MIO to optimise remaining parameters that only occur within the MIO algorithm.

The results of the optimisation process indicates that a codon range of [2, 20] works best, which means that a single test case may contain up to 20 events. Furthermore, a probability of 30% for applying extension and reduction local search indicates that the search benefits from using these operators. For the MOSA algorithm, we chose a population size of 30 and a crossover probability of 70%. Regarding MIO, an unreachable focus phase of 100% combined with a random test generation probability starting at 90% reveals that exploration is more beneficial than exploitation for the SCRATCH problem domain. A complete list of all used parameter configurations can be found in the WHISKER repository.

6.1.6 Environment

We conducted our experiments in a controlled execution environment using a Docker image based on Debian Slim Buster and Node.js 16. The used revision of WHISKER is identified by git commit 8537271. The experiments were run on a dedicated computing cluster. Each of its nodes features one Intel Xeon E5-2690v2 CPU with 3.00 GHz and 64 GB of RAM. Each run of WHISKER was allocated one CPU core and 5 GB of RAM.

6.2 Experiment Methodology

6.2.1 RQ1: How Much Can Test Execution be Accelerated Reliably?

To confirm whether the changes introduced to the SCRATCH VM allow for *reliable* accelerated test execution without introducing any form of non-deterministic behaviour, we run the manually written tests in the MANUAL dataset presented in Section 6.1.1 with a fixed seed for the random number generator, and repeat the experiment 20 times. If the tests reveal no functional differences in the execution, we can then validate whether the SCRATCH VM can be accelerated at all by comparing the total execution time on all 15 projects using varying acceleration factors. Finally, we want to ascertain that the accelerated execution of SCRATCH programs behaves deterministically in two ways: First, the test results of a given project for each used acceleration factor is compared with the outcomes of the non-accelerated test execution to make sure that the introduced acceleration of SCRATCH programs does not alter the execution behaviour. Second, to validate that the modified SCRATCH VM does not introduce any flakiness, we also check whether all 20 experiment repetitions within the observed acceleration factors lead to the same test results.

6.2.2 RQ2: Can SCRATCH Projects be Trivially Covered?

In order to determine whether SCRATCH projects represent a test generation challenge in the first place, we run a baseline algorithm of random testing with dynamic event selection on the Random1000 and Top1000 datasets. While there is no clear boundary of what constitutes a “trivial” project, intuitively a project does not pose a challenge to a test generator if it is possible to consistently achieve 100% code coverage without requiring a large number of test executions and without using an evolutionary algorithm. Following this intuition, we

define for both datasets a threshold, which is based on the lowest number of test executions at which we encounter the first program that is not covered entirely. Using this threshold, we then report the number of trivial and non-trivial projects for each dataset individually.

6.2.3 RQ3: What is the Best Test Generation Algorithm for SCRATCH Programs?

Section 5 describes three different algorithms for test generation; the aim of this research question is to determine which of these performs best. To answer this question, each search algorithm is applied 20 times to every project of the Random1000 and Top1000 dataset using the dynamic event selector and a search budget of 10 minutes. Besides the achieved block coverages, we also compare the number of events contained within final test suites and the average execution time of a test case during the search by comparing averages, the \hat{A}_{12} effect size and the number of projects for which one approach outperforms another one. Furthermore, we report the average block coverage achieved over time.

6.2.4 RQ4: How Effective are Generated Tests at Detecting Faults?










In our fourth research question, we evaluate whether the generated tests together with regression assertions (Section 4.3) are able to detect faulty SCRATCH programs. For that purpose, we conduct two experiments that seek to answer the research question from two different angles: First, we assess the generalisability of our approach by generating a large dataset of faulty programs based on a few selected program mutation operators. The second experiment then evaluates the applicability in a real-world scenario.

Within the first experiment, we use a mutation analysis framework that implements the eight mutation operators shown in Table 5. The selection of mutation operators is based on the traditional set of sufficient mutation operators (Offutt et al. 1996). Using the mutation framework and the tests generated during RQ3, we produce mutants for each program and validate whether the tests with assertions are able to detect the inserted program modifications. A mutant is detected if a test that passes on the original program leads to a failing assertion when executed on the mutant.

During the mutation analysis, we first load the test suite of the original program, before executing it against all its mutants one by one. This revealed a problem in the memory management of the SCRATCH VM: Every time a project is loaded, the VM deserializes associated assets (costumes and sounds), and stores them in memory. All mutants share the same assets, but previously loaded assets are neither reused nor cleared when loading the next mutant. Instead, a redundant copy of the data is created in memory. While this memory leak is not a problem for our test and assertion generation, where we reset the program state directly between test executions, it is a problem during mutation analysis, where not only the state but also the code need to be replaced. We therefore limit the RQ4 experiments to the Random1000 projects, which are small enough such that the memory leak does not affect the results significantly, as only 4% of all test executions escalated into program crashes. The Top1000 set, on the other hand, contains a rich set of assets, and we observed crashes in 58% of all test executions. The maintainers of the SCRATCH VM are aware of the memory leak,³ and while the problem is unresolved at the time of writing, we seek to evaluate the effectiveness of our automatically generated assertions on bigger programs (e.g., the Top1000 set) in future studies.

³<https://github.com/LLK/scratch-vm/issues/1955>

Table 5 Mutation operators

Operator	Description	Affected Blocks
Key Replacement Mutation (KRM)	Replaces the <i>key</i> argument with a randomly chosen one.	
Single Block Deletion (SBD)	Removes a block that is neither a hat nor a branching block.	All except  
Script Deletion Mutation (SDM)	Deletes a script by removing a hat block.	
Arithmetic Operator Replacement (AOR)	Replaces an arithmetic operator with a randomly chose one.	
Logical Operator Replacement (LOR)	Replaces a logical operator with the opposing one.	
Relational Operator Replacement (ROR)	Replaces a relational operator with a randomly chosen one.	
Negate Conditional Mutation (NCM)	Negates boolean blocks.	
Variable Replacement Mutation (VRM)	Replaces a variable with a randomly chosen one.	

The generalisability of our approach is evaluated by reporting the *Mutation Score* (Jia and Harman 2010) for every applied operator after excluding test cases which reported a false-positive result on the respective unmodified project. In order to keep the mutation analysis experiment within a reasonable time frame, we only consider first-order mutants as is usually done, and set a timeout of 90 minutes for the mutation analysis of an individual project. To avoid false-positive results due to randomised program behaviour, we seed each test execution with the same seed that was used during the test generation phase.

In our second experiment for RQ4, we evaluate the applicability of our approach in a real-world application scenario by first generating tests for each model solution of the collected student submissions from the BUGS dataset presented in Section 6.1.2. Then, we run the generated tests on the corresponding student-submitted programs and verify whether our tests come to the same Pass/Fail conclusions as the human examiner. Finally, the research question is answered by reporting the precision, recall and F1-score based on the goal of detecting faulty student submissions. To account for random influences during test generation, we repeat both experiments 20 times.

6.3 Threats to Validity

Threats to Internal Validity To ensure that results can be trusted, WHISKER has an extensive test suite, RQ1 aims to demonstrate validity, and we manually inspected results. Upstream changes to the SCRATCH VM may require adaptation of our modifications. However, such changes are rather unlikely, as they would also break many programs shared across the Scratch community. Since WHISKER uses randomised algorithms and results may be affected by chance, all experiments are based on 20 repetitions and are statistically analysed following common guidelines (Arcuri and Briand 2014). The performance of search algorithms depends on many parameters. In order to ensure our results are not negatively influenced by unsuitable configurations, we applied the tuning procedure described in Section 6.1.5.

Threats to External Validity Results may not generalise beyond the specific dataset used for experiments. However, we aimed to maximise generalisability by using two large datasets of 1000 projects each for RQ2 and RQ3, one randomly sampled from the SCRATCH website, and the other based on popularity as measured using the number of “love” reactions from other users. A possible source of bias is that only projects which users chose to publicly share can be accessed this way. It is conceivable that programs not shared publicly are more incomplete or broken. Similarly, the 15 projects and their tests used to answer RQ1 may not suffice to cover all possible sources of non-determinism that may occur in SCRATCH projects. The reported results in RQ4 are based on the Random1000 dataset and may not generalise to more complex programs, such as the ones contained in the Top1000 set.

Threats to Construct Validity The main metric for comparison is code coverage (coverage of statement blocks). Code coverage is the most common metric used in practice as well as in research in order to compare test suites as well as test generation algorithms; however, whether and how code coverage is related to fault detection is an ongoing debate (Inozemtseva and Holmes 2014; Chen et al. 2020). We also evaluate whether the tests are able to detect artificial faults using mutation analysis. Mutation scores may be skewed by equivalent mutants (Budd and Angluin 1982). Furthermore, artificially generated mutants may not be representative of real program faults (Gopinath et al. 2014). However, automated testing is not the only targeted application scenario of WHISKER; the tests are intended to enable any form of dynamic analysis that can support the generation of hints and feedback to learners. Code coverage is a prerequisite for any form of testing and dynamic analysis, and so it is important to consider this as a first step.

7 Results

7.1 RQ1: How Much Can Test Execution be Accelerated Reliably?

To determine whether SCRATCH projects can be accelerated at all, we executed the 15 test suites of the MANUAL dataset and analysed the required execution durations for each project. As can be observed in Fig. 12a, every recorded data point resides beneath the diagonal, indicating that all projects benefit from an increased acceleration factor. Figure 12b shows its direct impact, with a factor of two halving the average execution time across all projects. However, as acceleration factors further increase, the gains in speed up are diminished significantly. This is due to halting blocks enforcing lock durations, which start to reach the minimum duration of a single step in the VM. The error bars displayed in red indicate that the execution speed is very consistent and does not suffer from significant variation, which shows that with the given hardware environment the results are consistent. However, hardware with higher code execution speed would be capable of processing individual steps faster, thereby making even higher speedups feasible.

Besides validating whether SCRATCH programs can be accelerated at all, RQ1 seeks to ascertain that acceleration does neither introduce flakiness nor alter the program’s behaviour as a whole. To ensure deterministic behaviour, Fig. 12c compares the number of flaky tests obtained using the accelerated SCRATCH VM (SVM) introduced in Section 3.1 against the improved accelerated SCRATCH VM (SVM+) which contains safety measures for sources of flaky behaviour as described in Section 3.2. In contrast to SVM, the SVM⁺ does not show any signs of non-deterministic behaviour within 20 repetitions of the same test as well as between the execution of a given test in an accelerated and non-accelerated scenario. The

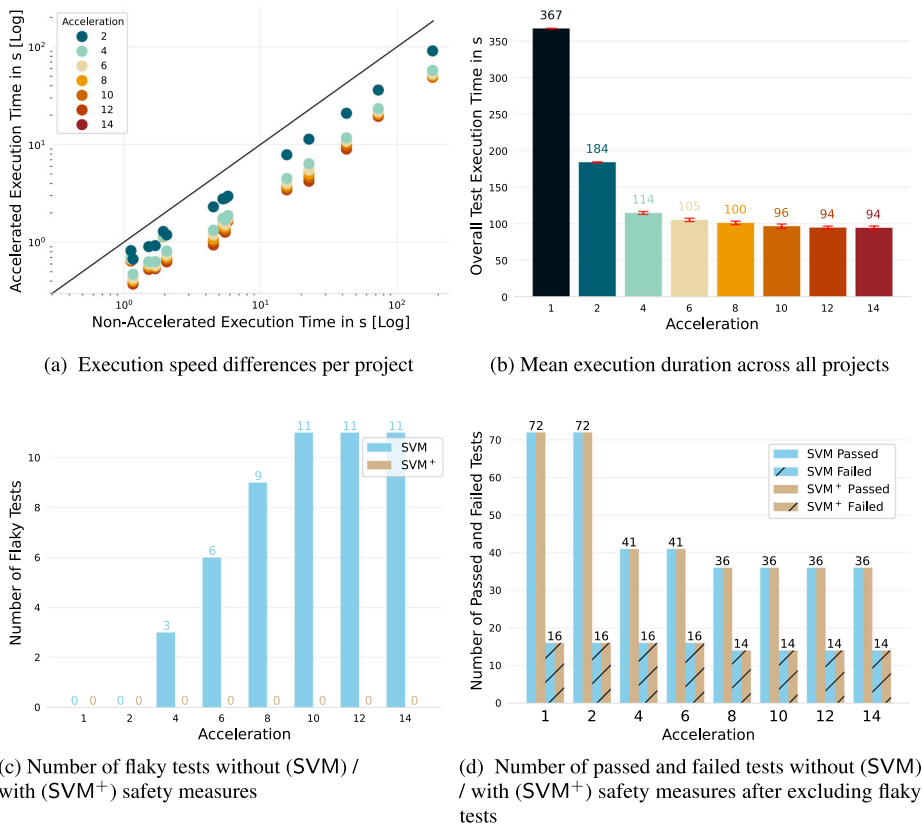


Fig. 12 Comparison of execution times, including standard deviations highlighted as red error bars, and flaky tests across different acceleration factors

SVM, however, shows increasingly flaky behaviour for higher acceleration factors due to reasons explained in Section 3.2. Even though the SVM does not show any signs of flakiness in our experimental setup for the unaccelerated scenario, non-deterministic behaviour is very likely to occur if tests are generated and executed on different machines or if the machine's computing resources are scarce during test execution.

Finally, Fig. 12d compares the number of passed and failed tests for both VM versions after excluding flaky (c.f. Section 3.2) test results. The results demonstrate that both variants achieve the same results, which indicates that the SVM⁺ behaves similarly to the SVM, and errors found in one of the two versions can be reproduced in the other one. All in all, due to the ensured determinism and limited increase in perceivable program speed up for acceleration factors greater than 10, we decided to conduct all following experiments using an acceleration factor of 10.

RQ1 Summary. Using an acceleration factor of 10, the overall execution duration for the test suite dataset was reduced by 73.8 % without changing the test results or introducing any flakiness. Thus, program execution can be sped up reliably if appropriate safety measures against sources of non-deterministic behaviour are added to the SCRATCH VM.

7.2 RQ2: Can SCRATCH Projects be Trivially Covered?

SCRATCH projects are created mainly by young learners, which raises the question whether test generation is actually a *problem*. For the Random1000 set, WHISKER managed to produce tests for 983/1000 projects across all used search algorithms. Most of the time, the search algorithms struggle with the remaining 17 programs due to memory and time limitations.

Figure 13a shows the distribution of coverage results using basic random testing on RANDOM1000, suggesting that the majority of projects are fully covered. Figure 14a visualises the relation of program size (measured in blocks), the number of executed tests until the time limit or 100% coverage was achieved, and the resulting average coverage. The plot contains two clusters: In the bottom left there is a large cluster of projects with less than 100 blocks, for which 100% coverage was achieved within less than 100 executed tests. The upper half contains a second cluster where results are more varied in terms of size, tests executed, and resulting coverage. Based on the project with the lowest average test execution count (plotted as ★ in Fig. 14a) for which the random test generator did not manage to reach full coverage, we define the RANDOM1000 dataset's threshold below which we classify projects as being trivial to be at 2.05 executed tests, leading to 391 trivial projects.

For TOP1000 the results look somewhat different (Fig. 13b): WHISKER manages to synthesise slightly fewer tests (947/1000) for the more challenging programs than for the Random1000 dataset. While there are still many projects covered fully, the coverage distribution shows a wider spread of coverage values achieved, with a tendency towards an almost bimodal distribution with one peak at about 100% and the other around 50%. Figure 14b confirms that the cluster of fully covered projects having less than 100 executed tests is very small for the popular projects. Furthermore, most projects have more than 100 statements, and many even more than 1000 blocks. This is considerable since the domain-specific blocks of the language do not require many blocks to conjure interesting behaviour. Furthermore, assembling the blocks for projects larger than 1000 statements in the SCRATCH code editor is a feat in itself. In the Top1000 dataset, we encounter the first project having a coverage below 100% (plotted as ★ in Fig. 14b) already after a test execution count of 1.50. Even though the difference between both thresholds is smaller than a single executed test, the Top1000 set contains significantly fewer trivial projects (29) than the Random1000 set.

Some of the trivial projects are simply very small. For example, Fig. 15 shows a project that contains two scripts with a total of ten statements, including two loops that control the




Fig. 13 Distribution of coverage results for random testing

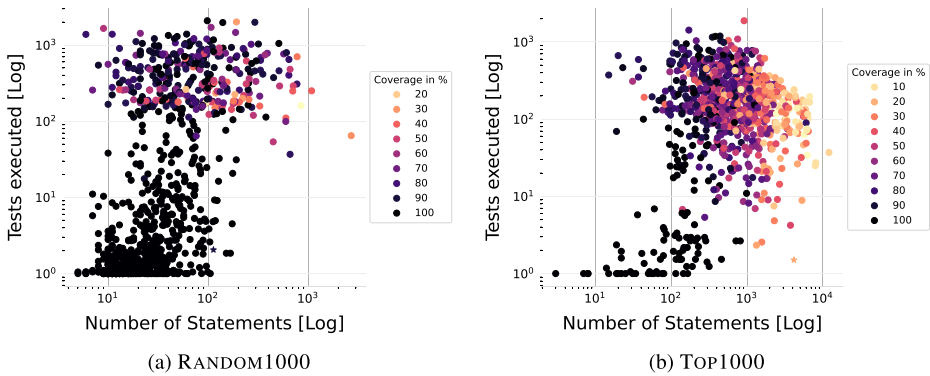


Fig. 14 Size vs. tests executed vs. coverage. Executed tests threshold after which projects are treated as non-trivial is marked as a star

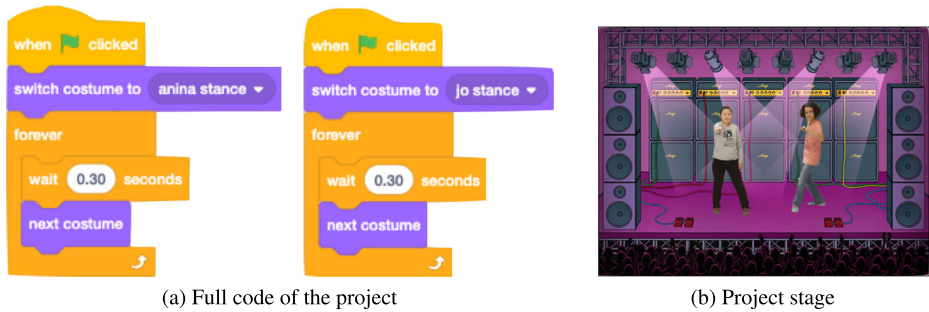


Fig. 15 Trivial example project (ID 400050176): Simply clicking on the green flag will cover everything within a few execution steps

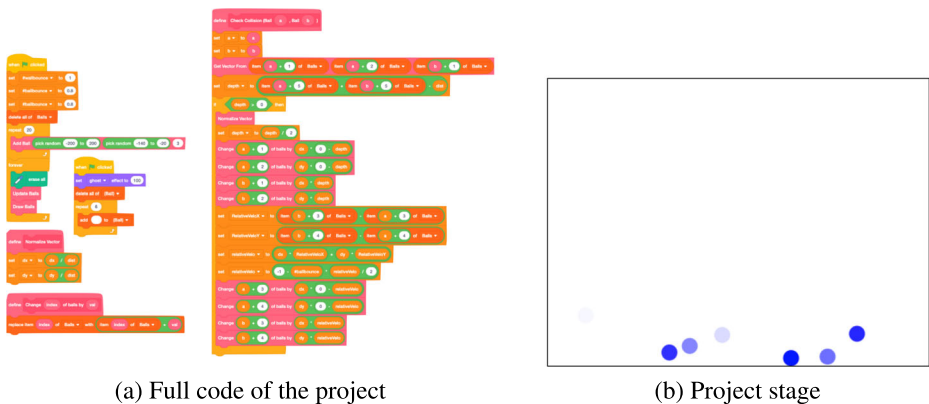


Fig. 16 Trivial example project (ID 400011212): Even though there are 314 blocks representing complex vector calculations, covering them requires no interactions




Fig. 17 Overall coverage

dance behaviour of the two sprites, which is achieved by cycling through costumes. Simply starting the program will cause all statements to be executed within a few execution steps. On the other hand, even projects that contain more code may be trivial if they are not interactive. For example, Fig. 16 shows a project which performs complex vector calculations on list datastructures, but this only serves to simulate bouncing balls with no user interactions.

Overall, the coverage observed on RANDOM1000 is clearly higher than in other domains, so it appears that many SCRATCH programs are indeed trivial. This matches the intended use case, where young learners initially take their first steps by building small animations and story-like projects. However, a fairly large number of projects nevertheless clearly challenges the random tester for projects in RANDOM1000, which suggests that even the average SCRATCH user may produce non-trivially covered projects. For projects to become popular (TOP1000) it rather seems that the level of difficulty is on par with other domains of software.

RQ2 Summary. Out of 983 randomly sampled projects, 391 achieve full coverage, requiring less than 2.05 randomly generated tests. For the 1000 most popular projects, only 29 can be covered fully, requiring less than 1.50 tests. The other 592 randomly sampled projects and most (918) of the Top1000 projects challenge random test generation.

7.3 RQ3: What is the Best Test Generation Algorithm for SCRATCH Programs?

The box plots in Fig. 17 show the overall coverage achieved by the different test generation algorithms. For Random1000 (Fig. 17a), all algorithms achieve very high coverage with a median of 100%. This is not surprising considering the large number of trivial projects (cf. RQ2). There is, however a difference noticeable on average, where random test generation leads to 92.7% coverage, MOSA achieves 95.4% coverage, and MIO achieves 95.5% coverage. For Top1000 (Fig. 17b) the coverage is substantially lower, and the differences between the algorithms are more pronounced: random test generation leads to 62.7% coverage, MOSA achieves 69.0% coverage, and MIO achieves 69.2% coverage.

The average coverage values suggest that MIO is the best algorithm, and MOSA is still better than random testing. Figure 18 sheds more light on the differences by showing the distribution of Vargha-Delaney \hat{A}_{12} effect sizes. The median is 0.5 for both datasets, which the statistical comparison summarised in Table 6 explains: For a large share of projects all

algorithms achieve the same level of coverage; for Random1000 this is often 100% (696 projects), while for Top1000 only 111 achieve 100% coverage. This means the chosen algorithm will very often make no difference, particularly for projects similar to Random1000.

The differences in average coverage can be explained by 198 projects in Random1000 where MIO performs better than random testing (significant for 159), and 174 for MOSA over random testing (significant for 146). This is substantially more than the number of cases where random is better than MIO (44, significant for 15) and MOSA (44, significant for 38). This is also reflected in the average effect size of 0.55 for MOSA vs. random testing, and 0.56 for MIO vs. random testing. Consequently, there are clear benefits to using either of the search algorithms over random testing on projects similar to those we randomly sampled.

The trade-off between search and random testing is less clear for Top1000: MIO performs better than random testing for 414 projects (significant for 306), and MOSA for 376 (significant for 293); at the same time, however, random is better than MIO for 376 projects (significant for 223) and MOSA for 406 (significant for 251). On average the effect size nevertheless leans towards search (0.55 for MIO vs. Random, and 0.53 for MOSA vs. Random). To better understand this result, Fig. 19 contrasts the coverage per project between random and the two search algorithms for both datasets. For all four cases the picture is very similar: A large share of the projects is clustered around the diagonal with equal coverage, and there is a larger spread of projects to the left of the diagonal than to the right, meaning that the search algorithms achieved higher coverage. For Top1000 (Fig. 20c and d) the spread around the diagonal is notably larger than for Random1000 (Fig. 20a and b), which shows that even though there are more cases with differences on Top1000, these are often insubstantial. When search is better, it is often better by a very large margin.

The differences between MOSA and MIO are small, but an \hat{A}_{12} of 0.51 for the Random1000 set and slightly more significantly better results for MIO (89 vs. 77) suggest that MIO overall is the algorithm better suited for the problem at hand. Figure 19 also shows only small differences between the two algorithms, confirming that differences are small, though slightly in favour of MIO. We conjecture that this is influenced by the larger degree of exploration achieved in MIO through the parameters that emerged from our tuning process.

Figure 20 shows how coverage evolves over time for both, Random1000 and Top1000: Random test generation very quickly converges at a lower coverage value on both datasets, whereas the two search algorithms successfully evolve tests to cover more code. The plot





Fig. 18 Effect sizes comparing algorithms wrt. coverage

Table 6 Test generation comparison of the achieved block coverages


Testset	Comparison	\hat{A}_{12}	Sig Better	Better	Equal	Worse	Sig Worse
Random1000	MOSA vs. Random	0.55	146	174	727	82	38
Random1000	MIO vs. Random	0.56	159	198	741	44	15
Random1000	MIO vs. MOSA	0.51	40	135	807	41	6
Top1000	MOSA vs. Random	0.53	293	376	165	406	251
Top1000	MIO vs. Random	0.55	306	414	157	376	223
Top1000	MIO vs. MOSA	0.50	89	342	278	327	77

Besides the \hat{A}_{12} value, we report the number of projects for which the first approach achieves significantly better, better, equal, worse and significantly worse results than the second approach. Statistically significant results are determined by using the Mann-Whitney-U test and a p -value < 0.05


also shows a distinct difference between MIO and MOSA: The MOSA algorithm requires longer to reach a higher level of coverage, whereas MIO has substantially higher coverage within the first 2–3 minutes of the search. This is due to the population based approach of MOSA, which applies evolutionary operators to entire generations of the population size chosen (30 in our experiments). In contrast, MIO produces one test at a time and directs the




(a) Random vs. MIO on Random1000



(b) Random vs. MOSA on Random1000



(c) Random vs. MIO on Top1000



(d) Random vs. MOSA on Top1000

Fig. 19 Comparison of achieved coverage

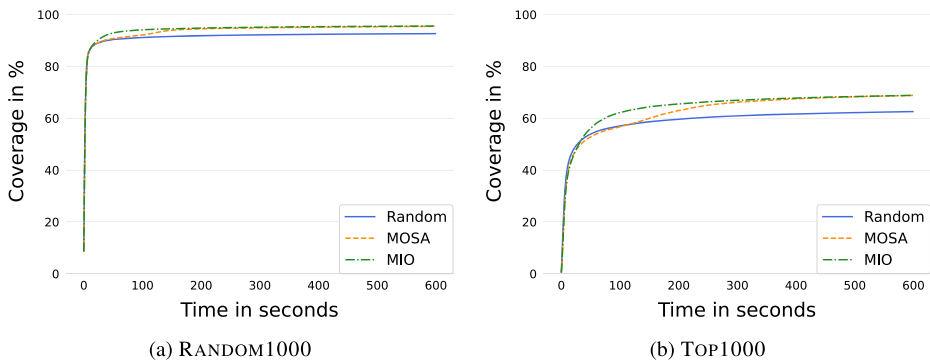


Fig. 20 Average coverage over time

search towards promising areas of the search space, which initially allows it to perform better. Consequently in particular if the time budget is limited, MIO may be a preferable choice.

The largest improvement of MOSA and MIO over random testing can be observed for projects that implement non-trivial story behaviour. For example, Fig. 21 shows project ID 401050644, where MOSA and MIO achieve 67.83 % and 66.24% coverage respectively, while random testing achieves only an average of 10.23%. The story consists of more than 100 individual scenes in which eight different sprites interact. Each scene is encoded as a script that is triggered by a message with the scene ID, and at the end broadcasts a message with the next scene ID. Covering the program entirely requires waiting long, and the fitness function provides a monotonic gradient to achieve this: The approach level captures the dependencies between the broadcasts, the control flow distance captures the progress in the scripts, and the branch distance captures the progression of time-related blocks. In conjunction with the extension local search, the problem thus becomes easy for the search.



(a) Code excerpt of the project



(b) Project stage

Fig. 21 Example project (ID 401050644): The project represents a story with more than 100 scenes, each encoded in an individual script, triggered by a broadcast

A similar pattern can be observed for many of the projects with large differences between search and random testing.

Although this type of project results in the largest difference between search and random testing, Fig. 22 suggests that very long execution sequences do not appear to be dominating. While the execution speed differences for the random set are negligible (Random: 2.81s, MOSA: 3.07s, MIO: 3.16s), we notice more significant differences and overall longer running tests for the Top1000 projects (Random: 9.14s, MOSA: 12.74s, MIO: 15.55s). However, a look at the execution times for projects of the Top1000 dataset in which all algorithms achieve exactly the same amount of coverage reveals that these results correlate with the increased program coverage (Random: 6.36s, MOSA: 6.78s, MIO: 7.07s), as the differences in execution times become negligible again.

Note that these execution times refer to *accelerated* execution, which means that effectively (unaccelerated) the tests are running up to an average of more than *two minutes* per test suite! Clearly, test generation without accelerated execution would be challenging. Notably these execution times are substantially higher than common values found in other test generation domains. For example, in search-based unit test generation (Fraser and Arcuri 2012) tests tend to execute within a few milliseconds. Since the computational costs of test execution are the central bottleneck in search-based test generation, this explains why, even though SCRATCH programs are substantially smaller than other types of software, we still have to run test generation for 10 minutes for reasonable results.

The search does not only provide advantages when the objective is to wait long enough. Figure 23 shows a game (ID 400148579) where the player controls the robot using the cursor keys, and the aim is to catch the star, which continuously moves to random positions. The script controlling the player score (Fig. 24a) provides a gradient in the fitness landscape that drives the search towards touching the star through the `touching Star` block, and the two if-conditions checking the score drive the search towards trying to repeat this. A further script checks for intermediate scores and displays messages. The search successfully controls the robot, and sometimes drags it, in order to reach scores that are substantially higher than those achieved by random testing. Indeed in most cases the search reaches the second (final) level of the game. Consequently, MOSA and MIO achieve an average of 92.57% and 92.71% coverage, respectively, whereas random testing only reaches an average of 80.14%.

While there are no projects where random testing achieves comparably large margins in terms of coverage over the search algorithms, there are some projects where random





Fig. 22 Average test execution time



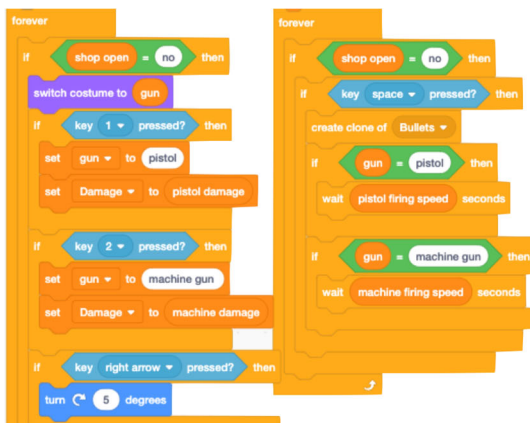
(a) Code excerpt




(b) Project stage

Fig. 23 Example project (ID 400148579): The user controls a robot that has to catch the star

testing does achieve higher coverage. Figure 24 shows a Zombie game where the player has to exterminate zombies without being eaten, using weapons that can be purchased in a shop. While the bullet-sprite provides some guidance for the search through a condition that checks if a zombie is touched, the fitness does not provide guidance towards shooting *all* zombies, nor to evade them. While random testing appears to be lucky nevertheless with an average coverage of 53.38%, MOSA only achieves an average of 49.68%. MIO benefits from the combination of exploration and exploitation and reaches an average coverage of 60.05%.



(a) Code excerpt of the project



(b) Project stage

Fig. 24 Example project (ID 402089829): Zombie shootout game

Many of the projects in the Top1000 dataset are games with similar challenges. For example, Fig. 25 shows “Paper Minecraft”, the most loved project of the Top1000 dataset. Containing 6720 statements, it is also among the biggest projects in the dataset (cf. Fig. 11b). Paper Minecraft implements a so-called sandbox game where players face no pre-determined objective but are encouraged to be creative by farming resources, creating buildings, etc. Random testing achieves an average coverage of 7.73%, while MOSA and MIO achieve 6.77% and 4.93%, respectively. To actually play the game (Fig. 25b) one has to select the “New Game” option on the title screen (Fig. 25a). Interestingly, when the button was hovered we observed a click-rate of 17/44 across all executions for random testing, compared to MOSA (7/22) and MIO (6/10). That is, random testing started the game more often than the other algorithms, which explains its higher coverage, while MIO started the game least often.

In general, for all algorithms the majority of tests for more complex games focus on interacting with the title screen rather than playing the actual game. Even when a test does play the game, the maximum length of 20 events per test case prevents it from doing so long enough, and only simple actions such as walking or switching items in the inventory can be performed in case of Paper Minecraft. Allowing longer tests would alleviate this problem to a certain extent, but would increase the computational costs; using variable length might also lead to effects of bloat (Fraser and Arcuri 2012). Like Paper Minecraft, many other games challenge the approach of optimising event-based sequences. Future work might consider reinforcement learning or other related techniques to address this problem.

Besides the achieved coverage, a further important aspect for consideration is the size of the resulting test suites, since these may need to be interpreted by users. As the tests generated by the algorithms vary in length, we quantify the length in terms of the overall number of events contained in a test suite. Since the size is influenced by the coverage achieved by a test suite (i.e., test suites with higher coverage tend to be larger), we compare the algorithms only on those projects, where all algorithms achieved the same coverage. Figure 26 summarises the average number of events in the final test suites for these projects, and Fig. 27 shows the distribution of effect sizes. The reported results are based on the test lengths *after* conducting the minimisation process, which reduces the average test suite size by an average of 39.12, 38.89 and 45.20 events on the Random1000 set and by 33.40, 34.10 and

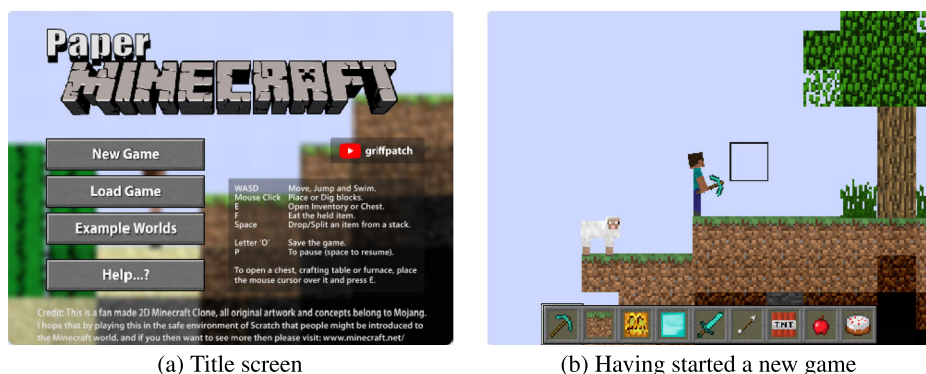


Fig. 25 Example project (ID 10128407): “Paper Minecraft” sandbox game

55.65 events on the Top1000 set for the Random, MOSA and MIO algorithm, respectively. This comparison shows for both datasets that MOSA and MIO produce smaller tests than random testing. The fact that this difference is measured *after* the minimisation suggests that the search algorithms succeed in finding targeted tests for more individual coverage goals, rather than accidentally covering many goals with long execution sequences. However, the minimisation has to remove the most events from MIO's test suites. We conjecture the longer event sequences to be influenced in particular by the extension local search, and how successful longer tests are replicated and mutated in MIO.

We use a test's JavaScript source to count its number of events. As every event selected by the test generator is followed by a Wait event for a single step, there are twice as many events as codons. Listing 2 shows an excerpt of a test for the game from Fig. 23 after removing automatically generated assertions. Events resulting from codons are followed by a single step (`t.runForSteps(1)`). KeyPress events are represented as two statements in the test code: The instruction to press the key for a certain number n of steps (`t.keyPress('.', n)`), and a Wait event with the same number n (`t.runForSteps(n)`). Figure 28 shows the average length of individual test cases for the different algorithms: Interestingly, for both datasets the median length of a single test covers between 49.73% and 78.90% of the entire suite, which indicates that most test suites consist of only 1–3 tests.

RQ3 Summary. MIO achieves the overall highest code coverage on SCRATCH projects in Random1000 with an average of 95.5 %, closely followed by MOSA with an average of 95.4 %, and random testing with an average of 92.7 %. On Top1000, the ranking is similar with 69.2 % (MIO), 69.0 % (MOSA), and 62.7 % (random testing).

7.4 RQ4: How Effective are Generated Tests at Detecting Faults?

For our last research question, we conducted two different experiments to evaluate our approach's effectiveness in generating test assertions that can detect faulty program behaviour. In our first experiment, we seek to assess the generalisability of our approach by executing the test suites produced in Section 7.3 on mutated versions of the programs and evaluate whether the synthesised tests are able to detect the generated mutants. Following common practice in mutation analysis and in order to ensure a sound evaluation, we exclude test cases which falsely mark a non-modified program as a mutant. However, as shown by Fig. 29, these false-positives only occur very rarely in the form of outliers and have a median




Fig. 26 Overall test suite size for projects with equal coverage




Fig. 27 Effect sizes comparing algorithms wrt. test suite size for projects with equal coverage

frequency of 0 %. The reason for these rarely occurring false-positives varies and is very program-specific. Out of 1273516 generated mutants, the tests generated by the random search algorithm were able to reach a *Mutation Score* of 52.23 %. In contrast, tests generated by MOSA and MIO detected 54.14 % and 55.67 % mutants from a total of 1250248 and 1253983 generated mutants, respectively. Please note that the total number of generated mutants varies slightly due to memory allocation issues, as explained in Section 6.2.4.

In addition to the frequency of false-positives, Fig. 29 also illustrates the distribution of killed mutants across the applied operators: For all three algorithms, the results are very similar, with MIO-generated tests detecting slightly more mutants per projects (50.28 %) than tests synthesised by MOSA (46.70 %) and random search (47.44 %). This small advantage originates from MIO's ability to achieve slightly higher coverage than the other algorithms, as shown in Section 7.3. All in all, the results demonstrate that the generated tests are able to detect faulty SCRATCH programs automatically. Nevertheless, further work should be done to improve the sensitivity of the generated assertions to reduce the frequency of false-positives and increase the number of detected program faults.

Finally, Fig. 29 reveals that certain mutants are harder to detect than others. NCM, KRM and SDM are easiest to detect since they can fundamentally alter program semantics, e.g., diverting control flow to the opposite branch in an if-else, or preventing the execution of entire scripts. In contrast, SBD, AOR and VRM show the lowest kill rates. Since they are

```
const test0 = async function (t) {
  t.dragSprite('Robot', 143, 126);
  await t.runForSteps(1);
  t.keyPress('right arrow', 3);
  await t.runForSteps(3);
  await t.runForSteps(1);
  t.keyPress('right arrow', 6);
  await t.runForSteps(6);
  await t.runForSteps(1);
  t.keyPress('up arrow', 1);
  await t.runForSteps(1);
  // ...
}
```

Listing 2 Test generated for project 400148579 (Fig. 23)




Fig. 28 Average length of test cases across all projects

applicable to many blocks, and not exclusively to the switching points of control flow, they have a lower chance of making impactful changes. We hypothesise that the former operators with their larger changes represent learners' mistakes well: Previous work (Frédérich et al. 2020) investigated typical bug patterns in the SCRATCH community, such as broadcast messages that are never sent or received, or cloning sprites without proper initialisation. These patterns are among the most common ones, and can be easily elicited by operators such as SDM or KRM, indicating that our mutants and tests can produce and detect common real-world faults in learners' programs. However, a closer analysis of fault coupling for SCRATCH mutation operators is out of scope for this paper and remains as future work.

The second experiment evaluates the effectiveness of our approach in the real-world application scenario of testing student submissions by comparing the results of tests that were generated on model solutions and executed on student submissions against the verdict of a human examiner. Since MIO proved to achieve slightly better results than Random and MOSA, we restrict this experiment to tests generated with MIO. To compute the precision, recall, and F1-score values, we define the goal of the experiment as detecting as many incorrect student submissions as possible. Figure 30 shows that MIO-generated tests achieve high F1-scores values for most program, indicating that the generated tests frequently come to the same conclusion as a human evaluator on most programs.




Fig. 29 Mutant kill rates across the applied mutation operators

For the programs *Garten*, *Geisterwald* and *Labyrinth*, the tests perform considerably worse than for other projects, thus reducing the overall mean F1-score to 0.63 per project. A closer look at Table 7 reveals the reasons for the low F1-scores in these three projects, which are the low precision values of 0.06, 0.19 and 0.18. Overall, the average precision and recall values of 0.59 and 0.81 indicate that our tests tend to be overly confident in marking submissions as incorrect. This behaviour originates from the assertion generation process, which adds assertions for nearly every property contained in the model solutions. As a consequence, the generated tests are very strict and only allow submissions that are very close or even identical to the model solution to pass the generated assertions.

In our experiment, deviations of the expected behaviour are primarily influenced by randomised behaviour of programs, as slight changes in the order of blocks can lead to a different consumption of the generated random numbers, resulting in small changes in the observed program properties. This behaviour is especially severe for the three projects *Garten*, *Geisterwald* and *Labyrinth*, as all three programs expose randomised program behaviour, and the number of correct programs is significantly higher than the number of incorrect ones. Although *Schatzinsel* and *Winter* have a similar distribution of correct and incorrect programs, precision is relatively high with values of 0.73 and 0.66 because these programs are less dependent on random number generators than *Garten*, *Geisterwald* and *Labyrinth*.

More generally, while overly strict tests help detect truly faulty programs, as shown by the relatively high recall value of 0.81, in an educational setting teachers may welcome alternative solution approaches that might lead to irrelevant assertion failures. Arguably, in this setting it would be reasonable to expect a teacher to select a subset of assertions or properties relevant to the assignment at hand. Alternatively, such issues may be overcome by minimising test assertions (Fraser and Zeller 2011), refining test assertions (Jahangirova et al. 2016), or exploring test oracles that capture the intended behaviour of programs with-




Fig. 30 F1-score of MIO-generated tests per project

Table 7 Number of correct/incorrect student submissions based on the verdict of the human evaluator together with the mean precision, recall and F1-score values of MIO-generated tests, given the goal of detecting as many incorrect student submissions as possible

Program	Correct	Incorrect	Precision	Recall	F1-score
Fangen	23	18	0.73	1.00	0.82
Garten	38	4	0.06	0.57	0.11
Geisterwald	34	8	0.19	0.12	0.15
HappyNewYear	29	13	0.80	1.00	0.89
Labyrinth	35	8	0.18	0.77	0.29
Papagei	26	17	0.84	0.93	0.88
Pong	17	24	0.58	0.85	0.69
Raumschiff	24	20	0.53	0.97	0.69
Schatzinsel	11	3	0.73	0.92	0.81
Skipiste	14	27	0.76	0.94	0.83
Snake	7	34	0.83	1.00	0.91
Sternenjagd	27	14	0.34	0.98	0.50
TicTacToe	17	25	1.00	0.34	0.50
Winter	36	6	0.66	1.00	0.80
Mean	24.14	15.79	0.59	0.81	0.63

out being disturbed by randomised program behaviour, such as model-based testing (Götz et al. 2022) or approaches based on artificial neural networks (Feldmeier and Fraser 2022).

Even though our assertions appear to be strict, the average recall of 0.81 lies slightly below the average coverage of 0.87. On the one hand, this may be because block coverage is a weak coverage criterion akin to statement coverage, and other coverage criteria may help revealing faults (Shamshiri et al. 2015). On the other hand, the projects *Geisterwald* and *TicTacToe* show particularly low recall values of 0.12 and 0.34. Such low recall values are often caused by the issue that blocks that would reveal the induced bugs are not reached by the tests. This issue may even occur for programs in which the tests reach generally high coverage values. For example, executed on the *TicTacToe* student submissions, the generated tests reach a high mean coverage of 90%. However, the execution of exactly these missed blocks would reveal the error that was induced to the model solution and not fixed by the students in incorrect submissions. Together with the first experiment in which MIO was able to achieve slightly better results due to reaching more block statements, it becomes apparent that high program coverage is crucial for detecting bugs in SCRATCH programs, and the high coverage values seen throughout all of our experiments should neither be interpreted to suggest testing SCRATCH programs is an easy nor a solved problem.

RQ4 Summary. The generated test suites detected more than half of the introduced program faults by achieving *Mutation Scores* of 52.23 %, 54.14 % and 55.67 % using Random Search, MIO and MOSA as test generation approaches, respectively. Lower precision than recall suggests that generated tests tend to be overly confident in classifying programs as faulty, but a high average F1-score of 0.63 shows that the generated assertions frequently come to the same conclusions as a human examiner.

8 Related Work

8.1 Automated Test Generation

A traditional approach to generate tests automatically is by using symbolic execution (Baldoni et al. 2018), which extracts path conditions from programs and then generates test inputs by solving the path conditions with constraint solvers. Symbolic execution is mostly used when testing at unit or API level, or when inputs can be represented as symbolic variables. In this paper, however, we consider system testing at the level of a user interface. While there have been attempts to apply symbolic execution also in this context (e.g., (Ganov et al. 2009; Mirzaei et al. 2012; Salvesen et al. 2015)), this is usually done to generate input values for specific user inputs (e.g., text). For SCRATCH programs, however, the challenge rather lies in finding timed sequences of simple user interactions. This problem is generally addressed using random and search-based test generation approaches.

Random testing of GUIs (Miller et al. 1995) consists of sending random user interactions to a program under test. Search-based testing generalizes this approach by adding objective functions, such as reaching target points in the source code (McMinn 2004), together with algorithms that optimize the inputs to reach the objectives. While the bulk of research on search-based testing considers function inputs or unit tests, the problem of generating tests for graphical user interfaces (GUIs) has also been successfully addressed using meta-heuristic search algorithms, for example in the context of Java Swing applications (Gross et al. 2012) or Android apps (Mao et al. 2016; Amalfitano et al. 2014; Mahmood et al. 2014).

Objective functions in search-based testing are usually based on the notion of code coverage, and require instrumentation to collect data that allows calculating fitness values. For some domains, such as Android apps, it is challenging to provide this instrumentation and to frequently execute long-running tests, therefore alternative black-box approaches have been proposed, e.g., aiming to maximise the amount of GUI changes observed (Mariani et al. 2012). In contrast, the size of SCRATCH programs represents no problems in terms of the scalability of fitness computations, and we therefore can base our fitness computations on inter-procedural analysis and instrumentation. However, test executions may still take a long time due to the time-based behaviour of SCRATCH programs.

An issue that is common to search-based GUI testing approaches is the difficulty of implementing search operators such as crossover, as crossing two sequences of events is likely to result in invalid sequences, where the events encoded in the sequences cannot be executed in the actual program states. Prior approaches to tackle this problem consisted of restricting crossover to suitable locations and ensuring valid sequences through repair (Mahmood et al. 2014), or using set-based representations where no sequences are modified during crossover (Mao et al. 2016). A common alternative is also to resort to heuristics that do not require such operators but, e.g., rather decide on executions based on probability distributions (Su et al. 2017). To overcome these problems of representation, we used an integer-based encoding based on grammatical evolution (O'Neill and Ryan 2001), which has not received much attention in the context of test generation yet (Anjum and Ryan 2020).

Which search algorithm is most effective is highly problem specific. Variants of random search are often sufficient (Shamshiri et al. 2018), but more advanced search algorithms provide clear benefits on more complex test problems. Our study confirms that this also holds in the domain of SCRATCH programs. At the unit testing level, it has been shown that searching for sets of tests that aim to cover all code (Fraser and Arcuri 2012) at once is most



effective (Campos et al. 2017; Panichella et al. 2018) using many-objective optimisation algorithms such as MOSA (Panichella et al. 2015) and MIO (Arcuri 2017), which is why we chose these many-objective optimisation algorithms also for our study.

8.2 Automated Testing and Analysis for SCRATCH Programs

Novice programming environments such as SCRATCH (Maloney et al. 2010) or SNAP (Harvey et al. 2013) are widely used in introductory programming curricula (Franklin et al. 2020; Garcia et al. 2015). These environments motivate students by enabling them to create programming artefacts that they can interact with, and they have been shown to improve learning gains and long-term interests towards programming (Weintrop and Wilensky 2017). Among the available programming environments, SCRATCH is by far the most popular environment, with the largest online youth programming community (Fields et al. 2017).

A core aspect of these programming environments is that they use blocks instead of text to avoid that learners have to memorise syntax or available programming commands. While this simplifies initial coding, students have been shown to still struggle in building logically coherent programs in SCRATCH (Meerbaum-Salant et al. 2011). They have also been reported to create “smelly” code (Aivaloglou and Hermans 2016; Hermans et al. 2016; Techapalokul and Tilevich 2017bb), and these code smells have been shown to have a negative impact on understanding (Hermans and Aivaloglou 2016). It is therefore important to provide tool-based support for learners as well as their teachers. The majority of prior work focused on statically analysing SCRATCH code. For example, the popular DR. SCRATCH (Moreno-León and Robles 2015) website assesses evidence of computational thinking in programs and can also point out code smells using the HAIRBALL (Boe et al. 2013) static analysis tool. Similar code smells are identified by QUALITY HOUND (Techapalokul and Tilevich 2017ba) and SAT (Chang et al. 2018). LITTERBOX (Frädrih et al. 2020) is an extensible framework that can identify not only code smells, but also patterns of common bugs as well as positive aspects such as code perfumes (Obermüller et al. 2021) in SCRATCH programs. A general limitation of these syntax-based approaches which we aim to address in this paper is that they can only provide limited reasoning about the actual and intended program behaviour.

Dynamic analysis is required to reason about program behaviour, and automated testing is a common means to enable dynamic analysis. Automated testing is commonly applied in the context of programming education for tasks such as assessing student programs to provide feedback after a task has been completed, or during its creation (Shute 2008). In many text-based programming environments, automated tests have been shown to enable various types of feedback, such as by displaying failed test cases (Edwards and Murali 2017), suggesting likely misconceptions (Gusukuma et al. 2018), and highlighting erroneous code (Edmison et al. 2017). Offering such immediate, automated feedback has been shown to improve students performance and learning outcomes (Corbett and Anderson 2001).

However, unlike text-based programming environments, novice programming environments like SCRATCH are often centered around custom graphical scenarios that are controlled by input streams of signals from users’ input devices such as keyboard and mouse, which causes challenges for automated testing. The ITCH tool (Johnson 2016) dynamically tests SCRATCH programs by translating a small subset of SCRATCH programs to Python code. However, such tests are limited to functions that take in static input/outputs, such as  and  blocks. Furthermore, ITCH does not automatically generate test cases. WHISKER (Stahlbauer et al. 2019) takes this approach a step further and, besides

execution of automated tests directly in SCRATCH, also provides automated property-based testing. SNAPCHECK (Wang et al. 2021b) applies similar concepts in the context of the SNAP! programming language. However, all of these existing testing tools focus on automatically executing manually written tests. In contrast, the aim of this paper is to automate the test generation process itself.

The work presented in this paper is integrated into WHISKER (Stahlbauer et al. 2019), but controls the SCRATCH VM directly and represents a separate component which is mainly connected with WHISKER through the result of the test generation, which is saved in WHISKER's format and can be re-executed with WHISKER. Our proof-of-concept on WHISKER test generation (Deiner et al. 2020) proposed a codon-based encoding, the use of interprocedural graphs to calculate fitness values, and accelerated test execution. This paper extends this initial proof-of-concept by providing an entirely new execution model, extending the codon encoding and search operators, providing new search algorithms, adding local search, refining the fitness function with the concept of control flow distance, adds many testability transformations to improve the fitness function, adds a new model for event extraction as well as new events, and adds test minimization as well as regression assertion generation, and adding many smaller technical improvements overall. In addition, a central contribution of this paper lies in the large empirical study.

9 Conclusions

The increasing popularity of block-based programming languages leads to a demand for tools to support programmers. However, even though languages like SCRATCH have millions of users, they lack fundamental analysis frameworks that are common for other programming languages, which inhibits the development of tools to support novice programmers. To address this issue, WHISKER makes it possible to run automated tests also on SCRATCH, but writing WHISKER tests remains challenging. In this paper we presented a fully automated approach to generate these tests given a SCRATCH program under test. Our experiments on three different, large datasets have demonstrated that automated test generation generally achieves very high coverage. This paves the way for advanced analysis and feedback tools.

Although our experiments suggest that WHISKER will fully cover many types of programs, we also observed two notable patterns of programs where the search-based test generation approach implemented by WHISKER could be improved:

- First, for many of the projects finding the correct sequence of user inputs is only part of the challenge, while in fact the *timing* is a more important question, and very often test generation would require waiting long durations for parts of the animations or sounds playing. While WHISKER accommodates for this through accelerated execution and including timing in the fitness function, the rather classical search-based testing approach that WHISKER implements nevertheless builds on the assumption that one can run many short executions. In contrast, many SCRATCH programs may be easier to test by alternative approaches aiming to drive individual, longer executions.
- Second, many SCRATCH programs, in particular the popular ones, represent games where a traditional test generation approach stands no chance of ever optimising a sequence that can really succeed at *playing* the game—which, alas, is a prerequisite to reaching interesting states and parts of the code. Possible avenues to address this problem will be to record and integrate user interactions with a program under test, or to apply reinforcement learning approaches to teach the computer to actually play the games.

The techniques described in this paper generalise conceptually in multiple dimensions: First, there are other block-based languages such as ALICE (Cooper et al. 2000) or SNAP (Harvey et al. 2013), which also use a similar concept of stages and sprites. Second, there are also text-based programming environments such as GREENFOOT (Kölling 2010) that are based on the same concept. Adapting our approach to these programming environments mainly requires engineering work to adapt our modified execution model, and to add support for different language constructs. We also anticipate that our deterministic execution model can influence the design of future programming environments. More generally, the encoding, search operators, and algorithmic modifications proposed in this paper are applicable in principle to any UI-based testing problem, independently of the underlying programming language.

Compared to other testing problems, the code coverage observed in our experiments is relatively high. Besides the smaller size of SCRATCH programs, one potentially influential factor is the absence of certain types of testing challenges such as external dependencies; for example, Android apps will frequently access web services and data storage, which leads to substantially lower code coverage (Mao et al. 2016). However, such challenges also exist in the domain of block-based languages: The SCRATCH language provides support for extensions that can provide arbitrary functionality, ranging from machine learning functionality to support for controlling external devices. Furthermore, there are related languages such as MBLOCK⁴, which extends SCRATCH with support for a wide range of robots. Supporting these features will require future work to extend our encoding as well as the underlying instrumentation.

Given the ability to generate tests for SCRATCH programs, we hope to enable new approaches for automated tutorial systems, automated repair systems, hint generation systems. To support this future work, WHISKER is available as open source at: <https://github.com/se2p/whisker>

Acknowledgements We thank Christoph Frädrieh, Sophia Geserer, and Niklas Zantner for their contributions to WHISKER. This work is supported by DFG project FR 2955/3-1 “TENDER-BLOCK: Testing, Debugging, and Repairing Blocks-based Programs”.

Funding Open Access funding enabled and organized by Projekt DEAL.

Data Availability A replication package containing all data from the experiments presented in this paper, as well as all artefacts and software needed to reproduce the results, is provided online at: <https://zenodo.org/record/7101776>.

Declarations

Conflict of Interests The authors declare that they have no conflict of interests.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

⁴<https://mblock.makeblock.com/>, last accessed September 2022

References

- Adler F, Fraser G, Gründinger E, Körber N, Labrenz S, Lerchenberger J, Lukasczyk S, Schweikl S (2021) Improving readability of SCRATCH programs with search-based refactoring. In: 21st IEEE international working conference on source code analysis and manipulation, SCAM, 2021. IEEE, pp 120–130
- Aivaloglou E, Hermans F (2016) How kids code and how we know: an exploratory study on the SCRATCH repository. In: Proceedings ICER, pp 53–61
- Amalfitano D, Fasolino AR, Tramontana P, Ta BD, Mobiguitar AMM (2014) Automated model-based testing of mobile apps. *IEEE Softw* 32(5):53–59
- Anjum MS, Ryan C (2020) Seeding grammars in grammatical evolution to improve search based software testing. In: European conference on genetic programming (part of EvoStar). Springer, pp 18–34
- Arcuri A (2013) It really does matter how you normalize the branch distance in search-based software testing. *Softw Test Verif Reliability* 23(2):119–147
- Arcuri A (2017) Many independent objective (MIO) algorithm for test suite generation. In: Proceedings of the international symposium on search based software engineering, SSBSE, vol 10452 of lecture notes in computer science. Springer, pp 3–17
- Arcuri A, Briand L (2014) A hitchhiker's guide to statistical tests for assessing randomized algorithms in software engineering. *Softw Testing Verification and Reliability* 24(3):219–250
- Baldoni R, Coppa E, D'elia DC, Demetrescu C, Finocchi I (2018) A survey of symbolic execution techniques. *ACM Comput Surveys (CSUR)* 51(3):1–39
- Boe B, Hill C, Len M, Dreschler C, Conrad P, Hairball DF (2013) Lint-inspired static analysis of SCRATCH projects. In: Proceedings SIGCSE, pp 215–220
- Budd TA, Angluin D (1982) Two notions of correctness and their relation to testing. *Acta Inform* 18(1):31–45
- Campos J, Ge Y, Fraser G, Eler M, Arcuri A (2017) An empirical evaluation of evolutionary algorithms for test suite generation. In: International symposium on search based software engineering. Springer, pp 33–48
- Chang Z, Sun Y, Wu T-Y, Guizani M (2018) Scratch analysis tool (sat): a modern SCRATCH project analysis tool based on antlr to assess computational thinking skills. In: 2018 14th International wireless communications & mobile computing conference (IWCMC). IEEE, pp 950–955
- Chen YT, Gopinath R, Tadakamalla A, Ernst MD, Holmes R, Fraser G, Ammann P, Just R (2020) Revisiting the relationship between fault detection, test adequacy criteria, and test set size. In: Proceedings of the 35th IEEE/ACM international conference on automated software engineering, pp 237–249
- Cooper S, Dann W, Pausch R, Pausch R (2000) Alice: a 3-d tool for introductory programming concepts. In: Journal of computing sciences in colleges. Consortium for computing sciences in colleges, vol 15, pp 107–116
- Corbett AT, Anderson JR (2001) Locus of feedback control in computer-based tutoring impact on learning rate, achievement and attitudes. In: Proceedings of the SIGCHI conference on human factors in computing systems. ACM, pp 245–252
- Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Trans Evol Computat* 6(2):182–197
- Deiner A, Frädich C, Fraser G, Geserer S, Zantner N (2020) Search-based testing for SCRATCH programs. In: International symposium on search based software engineering. Springer, pp 58–72
- Diner D, Fraser G, Schweikl S, Stahlbauer A (2021) Generating timed ui tests from counterexamples. In: International conference on tests and proofs. Springer, pp 53–71
- Edmison B, Edwards SH, Pérez-quíñones MA (2017) Using spectrum-based fault location and heatmaps to express debugging suggestions to student programmers. In: Proceedings of the nineteenth australasian computing education conference, ACE '17. Association for computing machinery, pp 48–54, New York
- Edwards SH, Murali KP (2017) Codeworkout: short programming exercises with built-in data collection. In: Proceedings of the 2017 ACM conference on innovation and technology in computer science education, pp 188–193
- Feldmeier P, Fraser G (2022) Neuroevolution-based generation of tests and oracles for games. In: 37th IEEE/ACM international conference on automated software engineering (ASE '22). ACM
- Fields DA, Kafai YB, Giang MT (2017) Youth computational participation in the wild understanding experience and equity in participating and programming in the online SCRATCH community. *ACM Trans Comput Educ (TOCE)* 17(3):1–22
- Frädich C, Obermüller F, Körber N, Heuer U, Fraser G (2020) Common bugs in SCRATCH programs. In: Proceedings of the 2020 ACM conference on innovation and technology in computer science education, pp 89–95

- Franklin D, Weintrop D, Palmer J, Coenraad M, Cobian M, Beck K, Rasmussen A, Krause S, White M, Anaya M, Crenshaw Z (2020) Scratch encore: the design and pilot of a culturally-relevant intermediate SCRATCH curriculum. In: Proceedings of the 51st ACM technical symposium on computer science education, SIGCSE '20. Association for computing machinery. ISBN 9781450367936, pp 794–800, New York
- Fraser G, Arcuri A (2012) Whole test suite generation. *IEEE Trans Softw Eng* 39(2):276–291
- Fraser G, Zeller A (2011) Mutation-driven generation of unit tests and oracles. *IEEE Trans Softw Eng* 38(2):278–292
- Fraser G, Arcuri A, McMinn P (2015) A memetic algorithm for whole test suite generation. *J Syst Softw* 103:311–327
- Fraser G, Heuer U, Körber N, Wasmeier E et al (2021) Litterbox: a linter for SCRATCH programs. In: 2021 IEEE/ACM 43rd international conference on software engineering: software engineering education and training (ICSE-SEET). IEEE, pp 183–188
- Ganov S, Killmar C, Khurshid S, Perry DE (2009) Event listener analysis and symbolic execution for testing gui applications. In: International conference on formal engineering methods. Springer, pp 69–87
- Garcia D, Harvey B, Barnes T (2015) The beauty and joy of computing. *ACM Inroads* 6(4):71–79
- Geldreich K, Funke A, Hubwieser P (2016) A programming circus for primary schools. In: ISSEP 2016, pp 49–50
- Gopinath R, Jensen C, Groce A (2014) Mutations: how close are they to real faults? In: 2014 IEEE 25th international symposium on software reliability engineering. IEEE, pp 189–200
- Götz K, Feldmeier P, Fraser G (2022) Model-based testing of SCRATCH programs. In: 2022 IEEE conference on software testing, verification and validation (ICST). IEEE, pp 411–421
- Greifenstein L, Obermüller F, Wasmeier E, Heuer U, Fraser G (2021) Effects of hints on debugging SCRATCH programs: an empirical study with primary school teachers in training. In: The 16th workshop in primary and secondary computing education, pp 1–10
- Gross F, Fraser G, Zeller A (2012) Search-based system testing: high coverage, no false alarms. In: Proceedings ISSTA, pp 67–77
- Gruber M, Lukaszczk S, Kroiß F, Fraser G (2021) An empirical study of flaky tests in python. In: 2021 14th IEEE conference on software testing, verification and validation (ICST). IEEE, pp 148–158
- Gusukuma L, Tech V, Cory Bart A, Kafura D, Ernst J (2018) Misconception-driven feedback : results from an experimental study, (1):160–168
- Harvey B, Garcia DD, Barnes T, Titterton N, Armendariz D, Segars L, Lemon E, Morris S, Paley J (2013) Snap!(build your own blocks). In: Proceedings of the 44th ACM technical symposium on computer science education, pp 759–759
- Hermans F, Aivaloglou E (2016) Do code smells hamper novice programming? a controlled experiment on SCRATCH programs. In: Proceedings ICPC. IEEE, pp 1–10
- Hermans F, Stolee KT, Hoepelman D (2016) Smells in block-based programming languages. In: Proceedings VL/HCC. IEEE, pp 68–72
- Inozemtseva L, Holmes R (2014) Coverage is not strongly correlated with test suite effectiveness. In: Proceedings of the 36th international conference on software engineering, pp 435–445
- Jahangirova G, Clark D, Harman M, Tonella P (2016) Test oracle assessment and improvement. In: Proceedings of the 25th international symposium on software testing and analysis, pp 247–258
- Jia Y, Harman M (2010) An analysis and survey of the development of mutation testing. *IEEE Trans Softw Eng* 37(5):649–678
- Johnson DE (2016) Itch: individual testing of computer homework for SCRATCH assignments. In: Proceedings SIGCSE, pp 223–227
- Kalliamvakou E, Gousios G, Blincoe K, Singer L, German DM, Damian D (2016) An in-depth study of the promises and perils of mining github. *Empir Softw Eng* 21(5):2035–2071
- Kölling M (2010) The greenfoot programming environment. *ACM Trans Comput Educ* 10(4):1–21
- Korel B (1990) Automated software test data generation. *IEEE Trans Softw Eng* 16(8):870–879
- Lee I, Martin F, Denner J, Coulter B, Allan W, Erickson J, Malyn-Smith J, Werner L (2011) Computational thinking for youth in practice. *Acem Inroads* 2(1):32–37
- Leitner A, Oriol M, Zeller A, Ciupa I, Meyer B (2007) Efficient unit test case minimization. In: Proceedings of the twenty-second IEEE/ACM international conference on automated software engineering, pp 417–420
- Luo Q, Hariri F, Eloussi L, Marinov D (2014) An empirical analysis of flaky tests. In: Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering, pp 643–653
- Mahmood R, Mirzaei N, Malek S (2014) Evodroid: segmented evolutionary testing of android apps. In: Proceedings of the 22nd ACM SIGSOFT international symposium on foundations of software engineering, pp 599–609

- Maj P, Siek K, Kovalenko A, Vitek J (2021) Codedj: Reproducible queries over large-scale software repositories. In: 35th European conference on object-oriented programming (ECOOP 2021). Schloss Dagstuhl-Leibniz-Zentrum für Informatik
- Maloney J, Resnick M, Rusk N, Silverman B, Eastmond E (2010) The SCRATCH programming language and environment. *TOCE* 10(4):1–15
- Mao K, Harman M, Jia Y (2016) Sapienz: multi-objective automated testing for android applications. In: Proceedings ISSTA, pp 94–105
- Mariani L, Pezze M, Riganelli O, Santoro M (2012) Autoblacktest: automatic black-box testing of interactive applications. In: 2012 IEEE fifth international conference on software testing, verification and validation. IEEE, pp 81–90
- McMinn P (2004) Search-based software test data generation: a survey. *Softw Test Verification Reliability* 14(2):105–156
- Meerbaum-Salant O, Armoni M, Ben-Ari M (2011) Habits of programming in SCRATCH. In: Proceedings of the 16th annual joint conference on Innovation and technology in computer science education, pp 168–172
- Miller BP, Koski D, Lee CP, Maganty V, Murthy R, Natarajan A, Steidl J (1995) Fuzz revisited: a re-examination of the reliability of unix utilities and services. Technical report, University of Wisconsin-Madison department of computer sciences
- Mirzaei N, Malek S, Păsăreanu CS, Esfahani N, Mahmood R (2012) Testing android apps through symbolic execution. *ACM SIGSOFT Softw Eng Notes* 37(6):1–5
- Moreno-León J, Robles G (2015) Dr. Scratch: a web tool to automatically evaluate SCRATCH projects. In: Proc. WIPSCCE, pp 132?–133
- Obermüller F, Bloch L, Greifenstein L, Heuer U, Fraser G (2021) Code perfumes: reporting good code to encourage learners. In: The 16th workshop in primary and secondary computing education, pp 1–10
- Offutt AJ, Lee A, Rothermel G, Untch RH, Zapf C (1996) An experimental determination of sufficient mutant operators. *ACM Transactions on Software Engineering and Methodology (TOSEM)* 5(2):99–118
- O'Neill M, Ryan C (2001) Grammatical evolution. *IEEE Trans Evol Comput* 5(4):349–358
- Panichella A, Kifetew FM, Tonella P (2015) Reformulating branch coverage as a many-objective optimization problem. In: 8th IEEE international conference on software testing, verification and validation (ICST), pp 1–?10
- Panichella A, Kifetew FM, Tonella P (2018) A large scale empirical comparison of state-of-the-art search-based test case generators. *Inf Softw Technol* 104:236–256
- Salvesen K, Galeotti JP, Gross F, Fraser G, Zeller A (2015) Using dynamic symbolic execution to generate inputs in search-based gui testing. In: 2015 IEEE/ACM 8th international workshop on search-based software testing. IEEE, pp 32–35
- Papert S (1980) *Mindstorms; Children Computers and Powerful Ideas*. Basic Book, New York
- Shamshiri S, Just R, Rojas JM, Fraser G, McMinn P, Arcuri A (2015) Do automatically generated unit tests find real faults? an empirical study of effectiveness and challenges (t). In: 2015 30th IEEE/ACM international conference on automated software engineering (ASE). IEEE, pp 201–211
- Shamshiri S, Rojas JM, Gazzola L, Fraser G, McMinn P, Mariani L, Arcuri A (2018) Random or evolutionary search for object-oriented test suite generation? *Softw Test Verification Reliability* 28(4):e1660
- Shute VJ (2008) Focus on formative feedback. *Rev Educ Res* 78(1):153–189
- Sirkkä T, Sorva J (2012) Exploring programming misconceptions: an analysis of student mistakes in visual program simulation exercises. In: Proceedings of the 12th Koli calling international conference on computing education research, pp 19–28
- Stahlbauer A, Kreis M, Fraser G (2019) Testing SCRATCH programs automatically. In: ESEC/SIGSOFT FSE. ACM, pp 165–175
- Su T, Meng G, Chen Y, Wu K, Yang W, Yao Y, Pu G, Liu Y, Su Z (2017) Guided, stochastic model-based gui testing of android apps. In: Proceedings of the 2017 11th joint meeting on foundations of software engineering, pp 245–256
- Techapolokul P, Tilevich E (2017b) Quality hound—an online code smell analyzer for SCRATCH programs. In: IEEE 2017 symposium on visual languages and human-centric computing (VL/HCC). IEEE, pp 337–338
- Techapolokul P, Tilevich E (2017b) Understanding recurring quality problems and their impact on code sharing in block-based software. In: Proceedings VL/HCC. IEEE, pp 43–51
- Vogl S, Schweikl S, Fraser G (2021) Encoding the certainty of boolean variables to improve the guidance for search-based test generation. In: Chicano F, Krawiec K (eds) GECCO '21: genetic and evolutionary computation conference. ACM, Lille, France, 10–14 Jul 2021, pp 1088–1096

- Wang W, Fraser G, Barnes T, Martens C, Price T (2021a) Execution-trace-based feature engineering to enable formative feedback on visual, interactive programs educational data mining in computer science education (CSEDM) workshop @ EDM'21
- Wang W, Zhang C, Stahlbauer A, Fraser G, Price T (2021b) Snapcheck: automated testing for snap programs. In: Proceedings of the 26th ACM conference on innovation and technology in computer science education, ITiCSE '21. ACM, pp 227–233
- Wegener J, Baresel A, Sthamer H (2001) Evolutionary test environment for automatic structural testing. *Inf Softw Technol* 43(14):841–854
- Weintrop D, Wilensky U (2017) Comparing block-based and text-based programming in high school computer science classrooms. *ACM Trans Comput Educ* 18(1):3
- Xie T (2006) Augmenting automatically generated unit-test suites with regression oracle checking. In: European conference on object-oriented programming. Springer, pp 380–403

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



Adina Deiner is a master student of Computer Science at the University of Passau. Her research interests include program analysis and programming education.



Patric Feldmeier is a PhD student of the Software Engineering II Chair at the University of Passau. His main research interests include search-based software engineering, software testing and reinforcement learning. Patric's primary research goal is to advance the state of fully automated Game and GUI testing using reinforcement learning approaches such as Neuroevolution. In addition, he is working on tools like Whisker that assist young learners throughout their developer journey by providing valuable feedback for the block-based programming language Scratch.



Gordon Fraser is a full professor in Computer Science at the University of Passau, Germany. He received a PhD in computer science from Graz University of Technology, Austria, worked as a post-doc at Saarland University, and was a Senior Lecturer at the University of Sheffield, UK. The central theme of his research is improving software quality, and his recent research concerns the prevention, detection, and removal of defects in software.



Sebastian Schweikl received the master's degree in computer science from the University of Passau, Germany, in 2019, where he is currently working as a PhD student. His main research interests include the design and analysis of search-based metaheuristics.



Wengran Wang is a PhD student of Computer Science at NC State University. Her research interests include building automated, personalized examples to support creative, open-ended programming.