



Software Engineering Group
Department of Computer Science
Nanjing University
<http://seg.nju.edu.cn>

Technical Report No. NJU-SEG-2018-IC-004

2018-IC-004

Genetic Synthesis of Concurrent Code Using Model Checking and Statistical Model Checking

Lei Bu, Doron Peled, Dachuan Shen, Yuan Zhuang

Model Checking Software - 25th International Symposium 2018

Most of the papers available from this document appear in print, and the corresponding copyright is held by the publisher. While the papers can be used for personal use, redistribution or reprinting for commercial purposes is prohibited.



Genetic Synthesis of Concurrent Code Using Model Checking and Statistical Model Checking

Lei Bu¹, Doron Peled²(✉), Dachuan Shen¹, and Yuan Zhuang¹

¹ State Key Laboratory for Novel Software Technology, Nanjing University, Nanjing, China

² Department of Computer Science, Bar Ilan University, Ramat Gan, Israel
doron.peled@gmail.com

Abstract. Genetic programming (GP) is a heuristic method for automatically generating code. It applies probabilistic-based generation and mutation of code, combined with “natural selection” principles, using a fitness function. Often, the fitness is calculated based on a large test suite. Recently, GP was applied for synthesizing correct-by-design concurrent code from temporal specification, where model checking was used for calculating the fitness function. A deficiency of this approach is that it uses a limited number of fitness values, based on a small number of modes for each verified specification property (e.g., satisfies, does not satisfy a given property). Furthermore, the need to apply model checking on many candidate solutions using the genetic process makes using an off-the-shelf model checker such as Spin prohibitively expensive. The repeated invocation of such a tool, compiling the code for a new candidate solution and running it, can render the performance of this approach several orders of magnitude slower than using an internal model checking. To tackle this problem, we describe here the use of a combination of statistical model checking, and a light use of model checking, for calculating the fitness required by GP.

1 Introduction

The classical approach for synthesis of interactive systems from temporal specification uses automata and game theory [25]. Synthesis of distributed or concurrent programs from temporal specifications is in general undecidable [26]. This calls for the use of heuristic methods. In particular, genetic programming based on model checking [15] employs a powerful heuristic search in the state space of candidate programs, which can be controlled and adjusted by an intelligent user.

The research in this paper was partially funded by an ISF-NSFC grant “Runtime Measuring and Checking of Cyber Physical Systems” (ISF award 2239/15, NSFC No. 61561146394). The authors from Nanjing University were also partially funded by a National Natural Science Foundation of China grant No. 61690204 and No. 61572249.

Genetic programming (GP) is an automatic method for generating code. It is based on beam search, i.e., a search that maintains in each generation a set of objects, rather than a *single* object. The search attempts to improve the quality of candidates from one generation to another, with mutual influence between candidates. Candidates propagate from one generation to the next one with probability based on their *fitness* value, which is an estimation on how close the candidate is from a correct solution. In addition, it uses the genetic operations of *mutation*, i.e., making small random changes to a candidate, and *crossover*, i.e., combining elements of two candidates. Since GP does not use backtracking, the only possibilities to deal with a failed search is to start with a new random seed or to try searching with a different fitness function.

Classical genetic programming is based on calculating the fitness function with respect to a large training set of test cases. Recently, using model checking for calculating the fitness function was studied in [14–18]. Model checking is a comprehensive approach for checking correctness, hence its use provides a greater assurance of the correctness of the code than testing. On the other hand, the number of correctness properties are typically not large, which results in a small set of fitness values. Additional fitness levels can be provided, e.g., “some executions satisfy a property”. However, even then, the fitness landscape is far from being smooth [8], which may sometimes limit the ability of the genetic search to converge.

Model checking based GP requires performing model checking for all the candidate programs generated during the process. Hence, it benefits greatly from having a dedicated model checker that is implemented within the genetic programming tool. Without it, the use of an off-the-shelf model checking tool like Spin [12] would be several orders of magnitude slower. This motivates using alternative ways of checking the fitness of candidates such as randomized testing and statistical model checking. However, these methods only provide a limited assurance about the correctness of the generated code.

In this work, we suggest to use statistical model checking (SMC) [20,27] to replace part of the work that is done by model checking. SMC is a simulation-based solution, which is less time and memory intensive than classical model checking. The goal of SMC is to check sample execution paths of the system and use methods like statistical hypothesis testing to calculate the probabilities of the system to satisfy a given property within a given statistical error. Compared with the model checking verdicts, probability measurement can provide smoother indication about how much the model satisfies a given property, which can assist in the convergence of the genetic process. Typical tools for SMC include Plasma [21] and Uppaal [3].

Applying SMC, which is based on finite executions, we immediately experience several inherent obstacles. The statistical sampling of the executions is limited to finite length, hence, the correctness of the generated programs is not guaranteed by the statistical evaluation. In particular, some properties may fail in very few executions (rare events), which may be missed in the statistical evaluation. Conversely, properties that hold for long or infinite execution sequences

may not be manifested during some of the finite executions that are checked. This suggests using a combination of SMC and model checking to achieve the best of both wofair scheds. Our approach combines SMC with light use of model checking, based on the Spin model checker, performed at the later stages of the genetic process.

2 Genetic Programming

During the 1970s, Holland [11] established the field known as *Genetic Algorithms* (GA). Individual candidate solutions are represented as fixed length strings of bits, corresponding to chromosomes in biological systems. Candidates are evaluated using a *fitness* function. The fitness approximates the distance of the candidate from a desired solution. Genetic algorithms evolves a *set of candidates* into a successor set. Each such set forms a *generation*, and there is no backtracking. Candidates are usually represented as fixed length strings. They progress from one generation to the next one according to one of the following cases:

- *Reproduction*. Part of the candidates are selected to propagate from one generation to the subsequent one. The reproduction is done at random, with probability relative to the relation between the fitness of the individual candidate and the average of fitness values in the current generation.
- *Crossover*. Some pairs of the candidates selected at random for reproduction are combined using the crossover operation. This operation takes parts of bit-strings from two parent solutions and combines them into two new solutions, which potentially inherit useful attributes from their parents.
- *Mutation*. This operation randomly alters the content of a small number of bits from candidates selected for reproduction (this can also be done after the crossover). One can decide on mutating each bit separately with some probability.

Unlike traditional point-by point search such as depth-first search or breadth-first search, the different candidates in a single generation have a combined effect on the search; progress tends to promote, improve and combine candidates that are better than others in the same generation. The process of selecting candidates from the previous generation and deciding whether to apply crossover or mutation continues until we complete a new generation. All generations are of some predefined fixed size N . This can be, typically, a number between 50 and 500. Genetic algorithms thus perform the following steps:

1. Randomly generate N initial candidates.
2. Evaluate the fitness of the candidates.
3. If a satisfactory solution is found, or the number of generations created exceeds a predefined limit (say hundreds or a few thousands), terminate.
4. Otherwise, select candidates for reproduction using randomization, proportional to the fitness values and apply crossover or mutation on some of them, again using randomization, until N candidates are obtained.
5. Go to step 2.

If the algorithm terminates unsuccessfully, we can restart it with a new random seed, or change the way we calculate the fitness function.

Genetic programming, suggested by Koza [19], is a direct successor of genetic algorithms. In GP, each individual organism represents a computer program. Programs are represented by variable length structures, such as trees (see Fig. 1) or a sequences of instructions. It is quite easy to transfer between a program and a syntax tree and vice versa. These trees are well typed. Each node is classified as *code*, *Boolean*, *condition* or *expression*. Leaf nodes are variables or constants, and other nodes have successors according to their type. For example, a *while* node (of type *code*) has one successor of type *Boolean* or *condition* and one successor of type *code* (for the loop body); a *Boolean* node and has two successors that can be of type *Boolean* or *condition*, and a *condition* node < has two successors of type *expression*. The genetic operations need to respect these (and possibly further) types, e.g., *expressions* cannot be exchanged with *Booleans*.

Crossover is performed by selecting subtrees on each of the parents, and then swapping between them. This forms two new syntax trees, having parts from both of their parents. There are several kinds of mutation operations. In *replacement* mutation, one picks at random a node in the tree, which is the root of a subtree. Then one throws away this subtree and replace it with a subtree of the same type, generated at random. In Fig. 1, the rightmost leaf node was chosen, which is marked with double ellipse. The subtree consists of this single node, representing the constant 1. Thus, it needs to be replaced with another expression, built at random. A new subtree was randomly generated, consisting of two nodes, representing the expression $a[0]$. In *insertion* mutation, a new node of the same type as the selected subtree is generated and is inserted just above it (type permitting); then one may need to complete the tree by constructing another descendant of the newly inserted node. For example, if we select an expression and insert above it a node that corresponds to addition $+$, it can be made one of the descendants to be summed up, say the left, but we need to complete the tree with a new right descendant. The *reduction* mutation has the opposite effect of *insertion*: the selected node is replaced with one of its offsprings (type permitting). In *deletion* mutation we remove the selected subtree, and recursively update the ancestors to make the program syntactically correct.

Syntax trees are not limited to a fixed size. Therefore the candidates can shrink or grow after mutation and crossover. In GP, there is actually a tendency of candidates to *bloat* with unnecessary code, for example, an assignment such as $a[1] := a[1]$. The countermeasure for this, called *parsimony pressure*, is to include a (small) negative value in the fitness function, corresponding to the length of the code. As a consequence, the resulting solutions are not expected to have a perfect fitness value, but instead they need to pass all the tests/verifications performed.

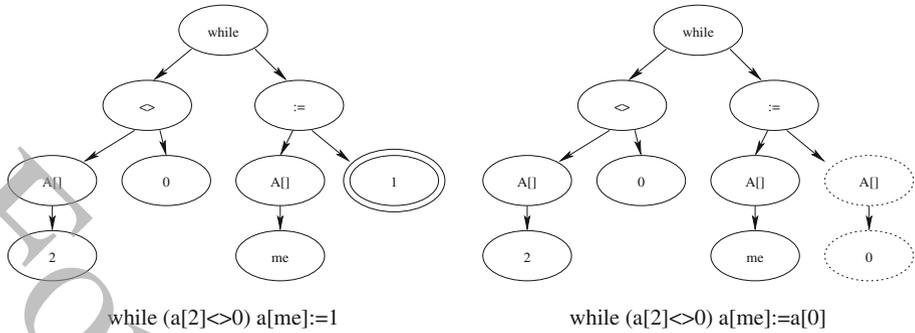


Fig. 1. Mutation on a syntax tree

3 GP Based on Model Checking and Statistical Model Checking

We want to employ GP to synthesize concurrent programs from given a temporal specification. We use linear temporal logic, LTL syntax [24]. The input includes, besides the temporal specification, also a *configuration*, which restricts the parameters of the desired solution. The configuration can restrict the depth of the generated syntax trees, the variable used, the allowed arithmetic and Boolean operators, and the number of processes. It can also contain a template that restricts the code, e.g., dictates that the code is embedded within a fixed loop or contains some fixed parts of code. The template has several uses:

- Using the template we can guarantee part of the behavior of the targeted code, simplifying the specification.
- We use as specification formalism LTL, which is limited to assert on *all* the executions of the code. On top of that, we can use the template to force checking different cases, providing some expressive power of branching temporal logic such as CTL [2]. Furthermore, templates also provide a testbed environment with *uncontrolled* actions, where the code needs to behave under all the interactions with it. This provides some expressive power of game logics [5].
- The template can be used to limit the state space of the search, e.g., suggesting that a solution will start with an assignment, or that it is embedded in a main loop. This can reduce the complexity of the genetic search and improve the chance and speed of coverage.

In [14–18], GP based on model checking was described and experimented with. The fitness function was solely calculated based on model checking results. Using model checking instead of testing to calculate the fitness function for GP allows a more reliable evidence of the correctness of the code. On the other hand, model checking is computationally expensive. In [15–18], it was observed that the number of specification properties is rather small, which creates a small number of fitness values. Therefore, a few intermediate levels were added on top of the

obvious *satisfies/does not satisfy* verdicts; in particular, levels such as *sometimes satisfies* and *satisfies with probability 1*.

3.1 A Running Example

As a running example, we look at synthesizing a solution for the well known *mutual exclusion* problem. Solutions for this problem from temporal specifications were synthesized using GP, where the fitness function is based on model checking [15, 16]. The configuration provided dictates the following structure:

<pre>p1: While W1 do nonCrit1:no-op preCS1 CS1:no-op postCS1 end while</pre>	<pre>p2: While W2 do nonCrit2:no-op preCS2 CS2:no-op postCS2 end while</pre>
--	--

The labels nonCrit_i represent the actions of the process p_i outside the critical section. The labels CS_i represent the critical section, which both processes want to enter a finite or unbounded number of times. These segments are not part of the synthesis task, and can be represented by trivial code `no-op`. The critical section is controlled by the code that will be synthesized for `preCS $_i$` and `postCS $_i$` . We require the following LTL properties:

Safety: $\Box \neg (p1 \text{ in } CS1 \wedge p1 \text{ in } CS2)$, i.e., there is no state where the program counters of both processes are in the critical section simultaneously.

Liveness: $\Box (p_i \text{ in } \text{preCS}_i \rightarrow \Diamond p_i \text{ in } \text{CS}_i)$, i.e., if a process wants to enter the critical section, then it will eventually succeed in doing so.

A solution that necessarily alternate between the process in entering their critical sections would also satisfy these conditions. Then, if one of them ceases to try entering its critical section, the other one can get blocked. To eliminate this problem, the variables W_i in the configuration are used to control whether processes want to keep entering the critical section. We can set different values to $W1$ and $W2$ to control the program behavior in different scenarios, including both processes want to enter the critical section, or only one wants to enter the critical section. This part of the code is fixed and not subject to synthesis.

Note that the configuration assures that the duration of the critical sections CS_i are finite. Hence there is no need to require that $\neg \Diamond \Box p_i \text{ in } \text{CS}_i$. We assume no (`goto`) statements are allowed, hence the synthesized parts are executed to completion each time they are entered.

3.2 Replacing Model Checking with Statistical Model Checking

Due to the two deficiencies of the use of model checking in genetic programming mentioned, complexity and lack of smoothness of the fitness value, we were motivated to replace part of its use by statistical model checking. In particular, we

generate for each GP candidate solution a large set of pseudo random executions; we check if these executions satisfy related specification properties.

The fitness function used in GP needs to be rather smooth in order to provide good convergence, and the statistical evaluation can provide multiple levels. Statistical evaluation may also be more affordable for some intricate synthesis problems. The simplicity of statistical methods is even further apparent for real time or cyber-physical systems. Another advantage that statistical model checking has over model checking is that it can be used for parametric systems and systems with infinite state space, where model checking has limited use for these applications.

For using statistical model checking over *finite* prefixes, we form a set of bounded temporal properties over *finite* prefixes of executions that are related to the original LTL properties over infinite sequences. Safety properties can be migrated directly to finite prefixes (A safety property is violated when there is a finite prefix that does so [1]). We use an additional temporal operator \blacklozenge , where $\blacklozenge\varphi$ holds in a sequence when φ holds at its last state. A finite prefix may present only partial information, and the property may be violated or satisfied only in a longer prefix. The properties over finite prefixes that correspond to the original properties *provide support* to the case that the original properties hold for the infinite sequences, but do not always guarantee them. For example, instead of the liveness property, we use a property that a process enters its critical section some fixed amount of times. The larger this number is, the more we are convinced that the liveness property holds. However, a large number will only be manifested in a long prefix. We pick up these related properties over finite prefixes according to our intuition (we may fine tune them if the genetic process fails).

At the moment we do not have a way of obtaining these related properties *automatically* from the original LTL properties, and this can be the subject of further research (e.g., using genetic co-evolution [19] or learning). Nevertheless, we do not expect the synthesis of concurrent programs to be completely automatic, as it was shown to be undecidable [26].

We illustrate the choice of related properties over finite prefixes for the running example. Suppose that we decide to check n executions, each one of them is limited to a length of k . We can fine tune the parameters n and k on several test runs to see what works. We can also try to estimate the size (number of states) of the desired solution to provide such parameters where errors will be found with high probability [10].

We split the original liveness property into several bounded properties. The predicates $enter_i$ represent the total number of times process p_i entered its critical section in the current prefix.

- B . The case that both processes want to enter the critical section. We enforce that by setting $(W1 \wedge W2)$. We add two counters $enter_1$ and $enter_2$ to indicate the times that each process enters the critical section. Out of which we have:
 - B_1 . Both processes succeed entering the critical section multiple times: $\rho_{B_1} = \blacklozenge(enter_1 > 1 \wedge enter_2 > 1)$.

- B_2 . One process enters the critical section multiple times and the other only once: $\rho_{B_2} = \blacklozenge((enter_1 > 1 \wedge enter_2 = 1) \vee (enter_1 = 1 \wedge enter_2 > 1))$.
- B_3 . Only one process succeeds in entering, or both enter exactly once: $\rho_{B_3} = \blacklozenge(enter_1 + enter_2 \geq 1 \wedge 0 \leq enter_1 \times enter_2 \leq 1)$.
- B_4 . Both processes do not succeed in entering their critical section $\rho_{B_4} = \blacklozenge(enter_1 + enter_2 = 0)$.
- O . Only one process p_1 wants to enter, when forcing $(W1 \wedge \neg W2)$. Out of which we have:
 - O_1 . The process p_1 succeeds entering its critical section multiple times: $\rho_{O_1} = \blacklozenge(enter_1 > 1)$
 - O_2 . The process p_1 succeeds entering its critical section only once: $\rho_{O_2} = \blacklozenge(enter_1 = 1)$
 - O_3 . The process p_1 does not succeed entering its critical section: $\rho_{O_3} = \blacklozenge(enter_1 = 0)$

We mark the SMC probabilities (as estimated by an SMC tool, or just the portions of executions satisfying each property among the randomly generated test cases) of the model satisfying these given properties by P_{B_1} , P_{B_2} , P_{B_3} , P_{B_4} , P_{O_1} , P_{O_2} and P_{O_3} respectively and the safety property by P_M . The fitness function is based on the above parameters.

We have the following coefficients, which can be assigned various values between 0 and 1:

- α multiplies P_M , the probability that the model satisfies the safety property.
- $\beta_1, \beta_2, \beta_3$ multiply P_{B_1} , P_{B_2} , P_{B_3} , the probability that the model satisfies ρ_{B_1} , ρ_{B_2} and ρ_{B_3} , respectively.
- γ_1, γ_2 multiply P_{O_1} , P_{O_2} the probability that the model satisfies ρ_{O_1} , and ρ_{O_2} , respectively¹.

We enforce that $\beta_3 < \beta_2 < \beta_1$, $\gamma_2 < \gamma_1$. A possible fitness function is

$$(\alpha \times P_M + \beta_1 \times P_{B_1} + \beta_2 \times P_{B_2} + \beta_3 \times P_{B_3} + \gamma_1 \times P_{O_1} + \gamma_2 \times P_{O_2}) \times 100$$

We normalize fitness to be between 0 and 100 by requiring that $\alpha + \beta_1 + \gamma_1 = 1$.

3.3 Problems and Solutions in Using SMC for Fitness Function

We need to pay attention to some issues in transforming SMC probabilities into fitness results. We will first list the difficulties, and then suggest some solutions.

Limited Distinction of the Probabilistic Approach. Although providing a smooth fitness value range, SMC based fitness function is only a rough estimate. In particular, it is hardly reasonable to assume that a solution that has 75% of its sampled prefixes satisfy some properties is uniformly worse than one in which 85% of the sampled prefixes satisfy them. However, the use of stochastic

¹ The coefficients for P_{B_4} and P_{O_3} are both 0, as these cases correspond to an inept solution.

selection of candidates for propagation by the genetic programming algorithm, where the given fitness only affects the *probability* of selecting the candidate, rather than directly selecting the best fitted ones, somewhat smoothens out the difference between such similar cases.

False Positives: Failure of Properties that Appear as Rare Events. The executions where an error is demonstrated may be rare; in which case one may need a lot of experiments and would, by chance, not catch the bad executions. For mutual exclusion, the processes *may* enter the critical section simultaneously, but on many executions they just independently enter and then exit, where the simultaneous stay within the critical section is not manifested on the selected random prefixes.

False Negatives: Negative Bias Due to Scheduling. Another problematic situation is where some liveness properties would not show up on a substantial number of prefixes due to scheduling. In a particular finite execution, a process may fail to enter the critical section since the other process is scheduled more frequently, although it could do so in a longer prefix or under a different scheduling.

Fairness. Many solutions of the mutual exclusion are based on some fairness assumption [23]; there, without allowing both processes ample opportunities to progress, the liveness will not hold. In particular, this is the case for the classical *Dekker* solution for mutual exclusion, presented by Dijkstra [4]. However, fairness is defined over infinite executions, and SMC checks only finite ones.

In order to tackle the above issues, which stem from the randomness and finiteness of the checked prefixes, we used a combination of the following ideas.

Extending the Measurements. Depending on the checked property, we may want to extend the measurements. For example, for the safety property, we can check more executions to increase the probability that we find the violation. For the liveness property, we may want to use longer executions to diminish the effect of unfair scheduling. These parameters are adjusted after some initial failures to synthesize correct solutions.

Using Combination of Cases. Because we cannot rely on fairness, and our tested sequences are finite, we should learn about the satisfaction of a property from observing the *combination* of the random checks. Take for example the case where we want to check that a process is not prohibited from entering the critical section. There may exist some prefixes where it fails to do so. However, if in a large percentage of the executions, it succeeds in entering the critical section, this can be used as evidence that the failure in the minority of the executions is due to unfavored scheduling. Then, given a certain threshold, we may apply “majority rules” to conclude that the liveness property holds. Accordingly, we may decide that when at least, say, 70% of the executions are satisfy ρ_{B_1} , the fitness treats all the executions in B as if they all satisfy ρ_{B_1} . Accordingly, in that case we use a simpler fitness function $(\alpha \times P_M + \beta_1 \times (P_{B_1} + P_{B_2} + P_{B_3}) + \gamma_1 \times P_{O_1} + \gamma_2 \times P_{O_2}) \times 100$.

However, this is not the only possible conclusion for this measurement: it may not be the discrimination of the scheduling that makes a process fail to enter

its critical section, but rather some scheduling that subsequently prevents the entrance to the critical section. Such a situation of multiple possible conclusions from the same statistical experiments can be resolved by the combined use of both majority rules *and* the light use of model checking (see below); model checking will catch such rare event errors that may otherwise not affect the fitness function.

Biasing the Probabilistic Selection. If we identify cases that may happen rarely, we can use biasing of the different choices in order to inspect them closer. For example, since catching violation of the safety property may be rare, we can reduce the probability of transitions that *exit* the critical section in favor of transitions of the other process that is outside the critical section. In essence, we are “waiting” for the other process to enter the critical section. For promoting liveness and providing more “fair” scheduling, we can decrease the probability of a transition of some process to be selected relative to the number of states where the other process has been waiting. A related technique for handling rare events, in the context of statistical model checking, is *importance splitting* [13], which split the test sequences into cases. Then one can zoom into checking cases where the rare events are believed to appear more frequently. This can also be a potential technique to handle the rare events problem.

Light Use of Model Checking as Certification or as Part of the Fitness. When candidates that receive very high fitness values are produced, in late generations, model checking can be used to certify that they indeed satisfy the desired properties. One can apply model checking sparingly, late on the genetic process, on candidates with already very high fitness value. We then may integrate the result of the model checking into the fitness and allow it to participate in additional generations.

Checking Ultimately Periodic Executions. We can replace checking finite executions by ultimately periodic ones. This can be done as in [7]. However, checking ultimately periodic sequences is more expensive than checking finite prefixes, as states on a sequence need to be hashed in order to detect cycles. This part was not implemented in our prototype.

4 Experiments

For each of the synthesis problems described above we performed experiments with SMC using Plasma [21]. Plasma uses Approximate Probabilistic Model Checking (APMC) [9] to provide a controlled accuracy on the statistical results². For accelerating the performance, we have also implemented an ad-hoc statistical evaluation algorithm which shares some of the merits of SMC. This implementation selects a given number of finite execution sequences and calculates the

² The configuration of running Plasma in our experiment includes the approximation threshold $\epsilon = 0.05$, and the confidence threshold $\delta = 0.01$. Please refer to [9] for detailed explanation of these parameters.

ratio of executions that satisfy a given property. However, it does not provide a *significance level* [6] for the measurement result.

The model checking is performed by *Spin*. *Spin* works here as separate software interfacing with ours, which needs to prepare its own (multiple) files and performs compilation on each candidate it checks, in order to make the verification; each activation of *Spin* by our code is slower than the statistical evaluations we make per candidate, hence we applied it sparingly. The *Spin* model checker was invoked when the fitness value reaches some threshold, which we set as 98. If model checking fails, we continue the genetic process, since the failed candidate solutions may still contain good “genetic material” so we can proceed from this point based on the SMC fitness calculations.

4.1 Synthesis of Solutions for Mutual Exclusion

The first set of experiments we conducted is to use GP to synthesize solutions of mutual exclusion. Without using *Spin* in the last stage to do the certification, our implementation can generate dozens of solutions that reach the highest fitness value easily. For example, three representative solutions (a), (b) and (c) are shown below. The processes are symmetric. The variables *me* and *other* can be concretized to i and $(i + 1) \bmod 2$ for process p_i ($1 \leq i \leq 2$) respectively.

While W1 do	While W1 do	While W1 do
v[me]=1	v[me]=1	v[me]=1
While (v[2]!=me) do	While (v[2]==other) do	While (v[other]!=0) do
v[2]=0	v[2]=1	While (v[other]==other) do
if(v[other]!=me)	if(v[other]!=other)	v[me]=0
v[2]=1	v[2]=0	end while
end while	end while	v[me]=1
CS	CS	end while
v[2]=other	v[2]=other	CS
end while	end while	v[me]=0
(a)	(b)	end while
		(c)

In the random simulations, both the processes show no violation of the mutual exclusion, starvation or deadlock. However, if we investigate the two solutions (a) and (b), we can find that they fail to satisfy the safety requirement. In some execution sequences, two processes can enter the critical section at the same time. Actually, for solution (a), among 10000 simulations, we could observe only 139 failures to satisfy the safety requirement. The unsafe scenario happened even fewer times in scenario (b): 4 times in 100000 simulations.

Solution (c) does not satisfy the liveness property. Actually, this solution represents the scenario where only if both processes want to enter the critical section indefinitely, then the liveness is satisfied; however, if one process decides to stop, then the other process will eventually be blocked forever. These examples demonstrate the problems raised in Sect. 3.3, where one may need a lot of

experiments and may, by chance, not catch the bad rare events. This leads us to the next experiment, where we used model checking as certification in the last generation of the genetic process.

For candidates that received fitness value of at least 98, we used the model checker Spin [12] to certify whether the desired properties are satisfied. To do that, we implemented an automatic generator to translate the solution generated by GP into the modeling language PROMELA of Spin. If the model checking confirmed correctness, the procedure was stopped. Otherwise, we continued the GP process (until the limit on the number of generations has been reached).

After integrating Spin to the GP process, we started to obtain correct solutions. One such solution is (*d*) below. This is a perfect solution that shares a similar structure with Dekker's algorithm. Another representative solution, (*e*), is similar to Peterson's algorithm. The difference between (*e*) and (*d*) is that (*e*) allows Boolean operators *and* and *or* in the conditions.

<pre> While W1 do v[me]=1 While (v[other]==1) do While (v[2]!=other) do v[me]=0 end while v[me]=1 end while CS v[me]=0 v[2]=me end while (<i>d</i>) </pre>	<pre> While W1 do v[me]=1 v[2]=me While((v[other]!=0) && (v[2]==me)) do end while CS v[2]=other v[me]=0 end while (<i>e</i>) </pre>
--	---

4.2 Synthesizing Solutions for Round Robin Scheduling

In this example, there are three processes p_0 , p_1 and p_2 , each with a critical section. The processes need to enter their critical sections in round robin order. That is, p_0 before p_1 , then p_2 , and repeating that order with p_0 , etc. The processes always want to enter the critical section (there is no flag W_i that restricts a process from wishing to enter). A trivial solution is that the processes would use a turn variable with three values, 0, 1 and 2, and each process will enter only when turn points to it and then increment it modulo 3. However, to make things less trivial, we require that we use only Boolean variables.

We allow solutions that are asymmetric in the sense that different values will be assigned in different processes. To allow that but still generate one candidate that will be concretized into three processes, we introduced to the generated process template as a syntactic construct an assignment statement of the form " $v[i]=b_0b_1b_2$ ", where $b_i \in \{0, 1\}$, $0 \leq i \leq 3$. This dictates that for the actual process p_i , the concretized statement will be $v[i]=b_i$. The variables which can show up in the solution are $v[0]$ to $v[3]$, and also $v[me]$, $v[other1]$, and $v[other2]$.

For each process p_i , me is concretized to i , $other1$ is concretized as $(i + 1) \bmod 3$, and $other2$ is concretized to $(i + 2) \bmod 3$.

Our GP based synthesizer generated several solutions similar to the following solution (f).

<pre> While true do v[3]=010 While(v[me]==001) do v[3]=101 v[3]=010 end while CS v[me]=001 v[other1]=101 end while (f) </pre>	<pre> While true do While(v[me]==001) do end while CS v[me]=001 v[other1]=101 end while (g) </pre>
---	--

We can see that each process in solution (f) refers to $v[me]$ and $v[other1]$. There are also assignments to $v[3]$ among the statements. However, as $v[3]$ is not used in any conditions at all, such statements can be safely removed. This is done here manually to demonstrate the solution, resulting in solution (g), and we did not implement parsimony pressure.

Let us concretize the solution of the three process p_0 , p_1 , and p_2 to (g_0) , (g_1) , and (g_2) respectively. Observe that only three variables are used in the solution, which makes this solution simple and elegant.

<pre> While true do While(v[0]==0) do end while CS v[0]=0 v[1]=1 end while (g0) </pre>	<pre> While true do While(v[1]==0) do end while CS v[1]=0 v[2]=0 end while (g1) </pre>	<pre> While true do While(v[2]==1) do end while CS v[2]=1 v[0]=1 end while (g2) </pre>
--	--	--

4.3 Synthesizing Solutions for Dining Philosopher

This synthesis problem involves several philosophers sit around a table; a philosopher can take the fork on his right or the one on his left as they become available. If a philosopher wants to eat, she must have both left and right forks and if he finishes eating, she needs to put the forks back and these forks will be available again. The problem is to design a concurrent algorithm with no deadlock and, under fair scheduling [22], no philosopher will starve.

To support this problem, we extend the basic variable library with semaphore variables, and also add semaphore-related operations such as *wait* and *signal* into the expression library. Two representative solutions generated by our method

are shown below in (h) and (i). It is interesting to see that although the GP synthesis processes allowed different kinds of programming constructs, including *while* loop, *if* condition and variable assignments in the expression library, the generated solutions are composed by only manipulation of semaphores. In these solutions, each philosopher (process) waits for the semaphore “mutex” to be free. Then she takes both the forks by capturing the semaphores “left” and “right”; these semaphores are translated to $i \bmod m$ and $i + 1 \bmod m$, for the i th philosopher. After finishes eating, a philosopher frees the semaphores she captured.

Both solutions (h) and (i) can pass the verification of Spin. In (h), when one philosopher is eating, all the other philosophers are blocked. Solution (i) permits more concurrency, as only philosophers that share forks with the dining ones will be blocked.

<pre> While true do think wait(mutex) wait(right) wait(left) eat signal(left) signal(right) signal(mutex) end while (h) </pre>	<pre> While true do think wait(mutex) wait(right) wait(left) signal(mutex) eat signal(right) signal(left) end while (i) </pre>
--	--

4.4 Performance Evaluation

We run the experiments on the SMC guided GP synthesis using Plasma, and using our own implementation of statistical evaluation. The corresponding data are marked with SMC and SE, respectively. The data for mutual exclusion, round robin and dining philosopher problems are marked as ME, RR, and DP, respectively.

In the experiment, we have 100 seeds for each generation, and if the GP does not generate a correct solution in 2000 generations, we abandon the current GP search. As the GP process involves randomness, we repeated the experiments for each problem 100 times. We record the average time each execution takes, the number of successful executions which generate perfect solution, the average number that the GP calls model checker per execution, and the average number of generations per execution. The performance data is given below in Table 1.

The built-in statistical evaluation (SE) algorithm basically implements a simplified method of SMC. We can see the success ratio of SE for each problem aligns with SMC, but performs almost two orders of magnitude faster. However, our built-in SE diminished the effect of having most of the overhead due to the invocation and repeated use of the tool. In fact, given that Plasma was used in this mode, its use should be considered very efficient. Moreover, the light need to

Table 1. Performance data

Problem	Method	Total executions	Average time	Success executions	Success rate	Average Spin call	Average generations
ME	SMC	21	7 h 53 m	0	0	280	2000
	SE	100	341.4 s	15	15%	378.3	1754
RH	SMC	22	7 h 25 m	8	36.4%	502	1751
	SE	100	479.4 s	41	41%	436.3	1389
DP	SMC	11	15 h 23 m	9	81.8%	1178	687
	SE	100	34 m 6 s	65	65%	1287	116

invoke Spin (in both modes), also as an external tool, did not make the entire synthesis process prohibitively expensive.

This experiments seem to support that our proposed SMC guided genetic synthesis method can be applicable and efficient in concurrent code synthesis with the help of a built-in implementation of SMC and model checking procedures. It can generate a correct solution with high success ratio within reasonable time overhead. Moreover, allowing SMC tools such as Plasma or Uppaal [3], and model checking tools such as Spin, to be integrated into a GP tool would make the genetic synthesis both efficient and powerful.

We do not compare our experiments directly to the results in [15–18] since there, an internal tool for performing model checking was used, which was tailored to provide additional levels besides the yes/no (+counterexample) that standard model checkers provide. But a powerful internal model checking is hard to implement, and can hardly compete with the breadth of a tool like Spin. Moreover, even with the additional model checking levels, the fitness function may not be smooth enough in some cases.

5 Conclusions

We described here the use of genetic programming based on statistical model checking for synthesizing concurrent code from its temporal specifications. Using statistical model checking for defining the fitness function has several advantages over using model checking. In particular, it can be more efficient, can be used in domains where model checking is not applicable, and can provide a smoother function, which helps to converge. We presented different ideas and parameters for defining statistical based fitness function.

We implemented these ideas and conducted experiments of synthesizing concurrent code. One of the main lessons we learned is that some common properties, such as mutual exclusion or eventual progress, may happen to be quite elusive in a model. This makes the tradeoff between efficiency and reliability. This also calls for using model checking to verify the generated solutions.

We used a hybrid approach, where we used statistical model checking for most of the duration of the genetic process, but involved model checking at the later part of the genetic process to certify the potential solution.

Our research on the combination of genetic programming, synthesis, statistical model checking and model checking already shows some encouraging results, but also calls for several follow ups. Besides improving our implementation (in particular, using built-in verification, as the connection with Spin and Plasma is quite time consuming) there are some interesting theoretical/practical directions. One direction is the use of biasing of the randomized experiments. Finally, we intend to make more experiments in synthesizing code, in particular of timed, probabilistic and cyber physical systems, where statistical approaches, such as statistical model checking, are found to be quite efficient.

References

1. Alpern, B., Schneider, F.B.: Recognizing safety and liveness. *Distrib. Comput.* **2**(3), 117–126 (1987)
2. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Kozen, D. (ed.) *Logic of Programs 1981*. LNCS, vol. 131, pp. 52–71. Springer, Heidelberg (1982). <https://doi.org/10.1007/BFb0025774>
3. David, A., Larsen, K.G., Legay, A., Mikucionis, M., Poulsen, D.B.: Uppaal SMC tutorial. *STTT* **17**(4), 397–415 (2015)
4. Dijkstra, E.W.: Cooperating sequential processes. In: Hansen, P.B. (ed.) *The Origin of Concurrent Programming*, pp. 65–138. Springer, New York (1968). https://doi.org/10.1007/978-1-4757-3472-0_2
5. Finkbeiner, B., Schewe, S.: Coordination logic. In: Dawar, A., Veith, H. (eds.) *CSL 2010*. LNCS, vol. 6247, pp. 305–319. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15205-4_25
6. Fisher, R.: *Statistical Methods for Research Workers*. Oliver and Boyd, Edinburgh (1925)
7. Grosu, R., Smolka, S.A.: Monte carlo model checking. In: Halbwachs, N., Zuck, L.D. (eds.) *TACAS 2005*. LNCS, vol. 3440, pp. 271–286. Springer, Heidelberg (2005). https://doi.org/10.1007/978-3-540-31980-1_18
8. Harman, M., Mansouri, S.A., Zhang, Y.: Search-based software engineering: trends, techniques and applications. *ACM Comput. Surv.* **45**(1), 11:1–11:61 (2012)
9. Hérault, T., Lassaigne, R., Magniette, F., Peyronnet, S.: Approximate probabilistic model checking. In: Steffen, B., Levi, G. (eds.) *VMCAI 2004*. LNCS, vol. 2937, pp. 73–84. Springer, Heidelberg (2004). https://doi.org/10.1007/978-3-540-24622-0_8
10. Hoeffding, W.: Probability inequalities for sums of bounded random variables. *J. Am. Stat. Assoc.* **58**(301), 13–30 (1963)
11. Holland, J.H.: *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control and Artificial Intelligence*. MIT Press, Cambridge (1992)
12. Holzmann, G.J.: *The SPIN Model Checker*. Pearson Education, Boston (2003)
13. Jegourel, C., Legay, A., Sedwards, S.: An effective heuristic for adaptive importance splitting in statistical model checking. In: Margaria, T., Steffen, B. (eds.) *ISoLA 2014*. LNCS, vol. 8803, pp. 143–159. Springer, Heidelberg (2014). https://doi.org/10.1007/978-3-662-45231-8_11
14. Johnson, C.G.: Genetic programming with fitness based on model checking. In: Ebner, M., O’Neill, M., Ekárt, A., Vanneschi, L., Esparcia-Alcázar, A.I. (eds.) *EuroGP 2007*. LNCS, vol. 4445, pp. 114–124. Springer, Heidelberg (2007). https://doi.org/10.1007/978-3-540-71605-1_11

15. Katz, G., Peled, D.: Genetic programming and model checking: synthesizing new mutual exclusion algorithms. In: Cha, S.S., Choi, J.-Y., Kim, M., Lee, I., Viswanathan, M. (eds.) ATVA 2008. LNCS, vol. 5311, pp. 33–47. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88387-6_5
16. Katz, G., Peled, D.: Model checking-based genetic programming with an application to mutual exclusion. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 141–156. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_11
17. Katz, G., Peled, D.: Synthesizing solutions to the leader election problem using model checking and genetic programming. In: Namjoshi, K., Zeller, A., Ziv, A. (eds.) HVC 2009. LNCS, vol. 6405, pp. 117–132. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-19237-1_13
18. Katz, G., Peled, D.: Code mutation in verification and automatic code correction. In: Esparza, J., Majumdar, R. (eds.) TACAS 2010. LNCS, vol. 6015, pp. 435–450. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-12002-2_36
19. Koza, J.R.: Genetic Programming: On the Programming of Computers by Means of Natural Selection. MIT Press, Cambridge (1992)
20. Legay, A., Delahaye, B., Bensalem, S.: Statistical model checking: an overview. In: Barringer, H., Falcone, Y., Finkbeiner, B., Havelund, K., Lee, I., Pace, G., Roşu, G., Sokolsky, O., Tillmann, N. (eds.) RV 2010. LNCS, vol. 6418, pp. 122–135. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-16612-9_11
21. Legay, A., Traonouez, L.-M.: Statistical model checking of simulink models with plasma lab. In: Artho, C., Ölveczky, P.C. (eds.) FTSCS 2015. CCIS, vol. 596, pp. 259–264. Springer, Cham (2016). https://doi.org/10.1007/978-3-319-29510-7_15
22. Lehmann, D.J., Rabin, M.O.: On the advantages of free choice: a symmetric and fully distributed solution to the dining philosophers problem. In: Conference Record of the Eighth Annual ACM Symposium on Principles of Programming Languages, Williamsburg, Virginia, USA, January 1981, pp. 133–138 (1981)
23. Manna, Z., Pnueli, A.: How to cook a temporal proof system for your pet language. In: Conference Record of the Tenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, January 1983, pp. 141–154 (1983)
24. Pnueli, A.: The temporal logic of programs. In: 18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October–1 November 1977, pp. 46–57 (1977)
25. Pnueli, A., Rosner, R.: On the synthesis of a reactive module. In: Conference Record of the Sixteenth Annual ACM Symposium on Principles of Programming Languages, Austin, Texas, USA, 11–13 January 1989, pp. 179–190 (1989)
26. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: 31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, 22–24 October 1990, vol. II, pp. 746–757 (1990)
27. Younes, H.L.S., Simmons, R.G.: Probabilistic verification of discrete event systems using acceptance sampling. In: Brinksma, E., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 223–235. Springer, Heidelberg (2002). https://doi.org/10.1007/3-540-45657-0_17