

Improving Swarming Using Genetic Algorithms

Etienne Renault

Abstract The verification of temporal properties against a given system may require the exploration of its full state space. In explicit model-checking this exploration uses a Depth-First-Search (DFS) and can be achieved with multiple randomized threads to increase performance.

Nonetheless the topology of the state-space and the exploration order can cap the speedup up to a certain number of threads. This paper proposes a new technique that aims to tackle this limitation by generating artificial initial states, using genetic algorithms. Threads are then launched from these states and thus explore different parts of the state space.

Our prototype implementation is 10% faster than state-of-the-art algorithms. These results demonstrate that this novel approach worth to be considered as a way to overcome existing limitations.

1 Introduction and Related Work

Model checking aims to check whether a system satisfies a property. Given a model of the system and a property, it explores all the possible configurations of the system, i.e., the *state space*, to check the validity of the property. Typically two kind of properties are distinguished, *safety* and *liveness* properties. This paper mainly focus on safety properties that are of special interest since they stipulate that some “bad thing” does not happen during execution. Nonetheless the adaptation of this work for checking liveness properties is detailed in Section 5.1.4.

The state-space exploration techniques for debugging and proving correctness of concurrent reactive systems has proven their efficiency during the last

decades [15, 20, 24, 4]. Nonetheless they suffer from the well known *state space explosion problem*, i.e., the state space can be far too large to be stored and thus explored in a reasonable time. This problem can be addressed using *symbolic* [5] or *explicit* techniques even if we only consider the latter one in this paper.

Many improvements have been proposed for explicit techniques. *On-the-fly exploration* [6] computes the successors of a state only when required by the algorithm. As a consequence, if the property does not hold, only a subset of the state space is constructed. *Partial Order Reductions (POR)* [26, 21, 17] avoid the systematic exploration of the state space by exploiting the interleaving semantic of concurrent systems. *State Space Caching* [11] saves memory by “forgetting” states that have already been visited causing the exploration to possibly revisit a state several times. *Bit-state Hashing* [13] is a semi-decision procedure in which each state is associated to a hash value. When two states share the same hash value, one of this two states (and thus its successors) will be ignored.

These techniques focus on reducing the memory footprint during the state-space exploration. Combining these techniques with modern computer architectures, i.e., many-core CPUs and large RAM memories, tends to shift from a memory problem to an execution time problem which is: *How this exploration can be achieved in a reasonable time?*

To address this issue multi-threaded (as well as distributed) exploration algorithms (that can be combined with previous techniques) have been developed [14, 3, 9, 20]. Most of these techniques rely on the *swarming* technique presented by Holzmann et al. [15]. In this approach, each thread runs an instance of a verification procedure but explores the state space with its own transition order.

Nowadays, best performance is obtained when combining swarming with *DFS*-based (Depth-First Search) verification¹ procedures [24, 4]. In these combinations, threads share information about states that have been *fully explored*, i.e. states where all successors have been visited by a thread. Such states are called *closed states*. These states are then avoided by other threads explorations since they can not participate in invalidating the property. These swarmed-DFS algorithms are linear but their scalability depends on two factors:

Topology problems. If the state space is linear (only one initial state, one successor per state), using more than one thread cannot achieve any speedup. This issue can be generalized to any state space that is deep but not wide.

Exploration order problems. States are tagged as *closed* following the DFS postorder of a thread. Thus, a state s can only be marked as *closed* after visiting at least N states, where N is the minimal distance between the initial state and s .

Table 1 highlights this scalability problem over the benchmark² used in this paper. It presents the cumulated exploration time (in a swarmed DFS fashion) for 38 models extracted from the literature. It can be observed that this algorithm achieves reasonable speedup up to 4 threads but is disappointing for 8 threads and 12 threads (the maximum we can test).

This paper proposes a novel technique that aims to keep improving the speedup as the number of threads increases and which is compatible with all memory reduction methods presented so far.

The basic idea is to use genetic algorithms to generate artificial initial states (Sections 2 and 3). Threads are then launched with their own verification procedure from these artificial states (Sections 4 and 5). We expect that these threads will explore parts of the state space that are relatively deep regarding to (many) DFS order(s). Thus, some states are tagged as *closed* without processing some path between the original initial state to these states.

Our prototype implementation (Section 6) has encouraging performances: the proposed approach runs 10% faster (with 12 threads) than state-of-the-art algorithms (with 12 threads). These results are encouraging and show that this novel approach worth to be considered as a way to overcome existing limitations.

¹ It should be noted that even if DFS-based algorithms are hard to parallelize [22] they scale better in practice than parallelized Breadth-First Search (BFS) algorithms.

² See Section 6 for more details about the benchmark.

This paper is an extension of our work published at VECOS'18 [23] where we proposed new parallel exploration algorithms built upon the generation of artificial initial states using genetic programming. These artificial states were then used on-the-fly to generate new seeds for the various threads used during the exploration. In this approach half of the threads were spawned from artificial states while the others used the classical approach used in model-checking algorithms. To handle the verification of safety properties, this approach was adapted to avoid (1) early termination and (2) reporting false positive.

In addition to the above (common with our previous paper [23]), we investigate one variant: the number of threads using artificial states may have an impact on the performances of our algorithms (more details in Section 6). We also detail the full proof for Algorithm 6 in Section 5.1.3 (while only its sketch was given in Renault [23]). Section 5.1.2 provides a detailed example of our algorithm while Section 5.1.4 focuses on reporting counterexample from our algorithms. Finally, Section 5.2 details how our algorithms can be combined with classical algorithms for checking liveness properties.

Related Work. To our knowledge, the combination of parallel state space exploration algorithms with the generation of artificial initial states using genetic algorithms has never been done. The closest work is probably the one of Godefroid and Khurshid [10] that suggests to use genetic programming as an heuristic to help random walks to select the *best* successor to explore. The generation of other initial states have been proposed to maximize the coverage of random walks [25]: to achieve this, a bounded BFS is performed to obtain a pool of states that can be used as seed states. This approach does not help the scalability when the average number of successors is quite low (typically when mixing with POR).

In the literature there are some work that combine model checking with genetic programming but they are not related to the work presented here: Katz and Peled [16] use it to synthesize parametric programs, while all the other approaches are based on the work of Ammann et al. [1] and focus on the automatic generation of *mutants* that can be seen as particular “test cases”.

2 Parallel State Space Exploration

Preliminaries. Concurrent reactive systems can be represented using *Transitions Systems* (TS). Such a system $T = \langle Q, \iota, \delta, V, \gamma \rangle$ is composed of a finite set of

	1 thread	2 threads	4 threads	8 threads	12 threads
Time in milliseconds	2960296	1796418	1186344	981222	978711
Speedup	1	1.65	2.50	3.016	3.025

Table 1 Problem statement about swarmed DFS like approaches.

states Q , an initial state $\iota \in Q$, a transition relation $\delta \subseteq Q \times Q$, a finite set of integer variables V and $\gamma : Q \rightarrow \mathbb{N}^{|V|}$ a function that associates to each state a unique assignment of all variables in V . For a state $s \in Q$, we denote by $\text{post}(s) = \{d \in Q \mid (s, d) \in \delta\}$ the set of its direct successors. A *path* of length $n \geq 1$ between two states $q, q' \in Q$ is a finite sequence of transitions $\rho = (s_1, d_1) \dots (s_n, d_n)$ with $s_1 = q$, $d_i = q'$, and $\forall i \in \{1, \dots, n-1\}$, $d_i = s_{i+1}$. A state q is *reachable* if there exists a path from the initial state ι to q .

Swarming. Checking temporal properties involves the exploration of (all or some part of) the state space of the system. Nowadays, best performance is obtained by combining on-the-fly exploration with parallel DFS reachability algorithms. Algorithm 1 presents such an algorithm.

This algorithm is presented recursively for the sake of clarity. Lines 4 and 5 represent the main procedure: `ParDFS` takes two parameters, the transition system and the number n of threads to use for the exploration. Line 5 only launches n instances of the procedure `DFS`. This last procedure takes three parameters, s the state to process, tid the current thread number and a *color* used to tag new visited states. Procedure `DFS` represents the core of the exploration. This exploration relies on a shared hashmap *visited* (defined line 2) that stores all states discovered so far by all threads and associate each state with a color (line 1):

- OPEN indicates that the state (or some of its successors) is currently processed by (at least) a thread,
- CLOSED indicates that the states and all its successors (direct or not) have been visited by some thread.

The `DFS` function starts (lines 7 to 8) by checking if the parameter s has already been inserted, by this thread or another one, in the *visited* map (line 7). If not, the state is inserted with the color OPEN (line 7). Otherwise, if s has already been inserted we have to check whether this state has been tagged CLOSED. In this case, s and all its successors have been visited: there is no need to revisit them. Line 10 grabs all the successors of the state s that are then shuffled to implement the swarming. Finally lines 11 to 15 perform the recursive DFS: for each successor s' of the current state, if s' has not been tagged CLOSED a recursive call

Algorithm 1: Parallel DFS Exploration.

```

1 enum color = { OPEN, CLOSED }
2 visited: hashmap of (Q, color) // Shared variable
3 stop ← ⊥ // Shared variable
4 Procedure ParDFS((Q, ι, δ, V, γ) : TS, n : Integer)
5   DFS(ι, 1, OPEN) || ... || DFS(ι, n, OPEN)
6 Procedure DFS(s ∈ Q, tid : Integer, status : color)
7   if s ∉ visited then visited.add(s, status)
8   else if visited[s] = CLOSED then return
9   // Shuffle successors using tid as seed
10  todo ← shuffle(post(s), tid)
11  while (¬stop ∧ ¬todo.isempty()) do
12    s' ← todo.pick()
13    if s' is in the current recursive DFS stack then
14      continue
15    if (s' ∉ visited) ∨ visited[s'] ≠ CLOSED then
16      DFS(s', tid, status)
17  visited[s] ← CLOSED
18  if (s = ι) then stop ← ⊤

```

is launched. When all successors have been visited, s can be marked CLOSED.

One can note that a shared Boolean *stop* is used in order to stop all threads as soon as a thread closes the initial state. This Boolean is useless for this algorithm since, when the first threads ends, all reachable states are tagged CLOSED and every thread is forced to backtrack. Nonetheless this Boolean will be useful later (see Section 4). Moreover the *visited* map is thread safe (and lock-free) so that it does not degrade performances of the algorithm.

Problem statement. The previous algorithm (or some adaptations of it [24, 4]) obtains the best performance for explicit model checking. Nonetheless this swarmed algorithm suffers from a scalability problem. Figure 1 describes a case where augmenting the threads will not bring any speedup³. This figure describes a transition system that is linear. The dotted transitions represent long paths of transitions. In this example, state x cannot be tagged CLOSED before state y and all the states between x and y have been tagged CLOSED. The problem here is that all threads start from state s . Since threads have similar throughput they will discover x and y approximately at the same time. Thus they cannot benefit from the information computed

³ This particular case will certainly degrade performance due to contention over the shared hashmap.

by the other threads. This example is pathological but can be generalized to any state space that is deep and narrow.

Suppose now that there are 2 threads and that the distance between s and x is the same than the distance between x and y . The only way to obtain the maximum speedup is to launch one thread with a DFS starting from s and launch the other thread from x . In this case, when the first thread reaches state x , x has just been tagged CLOSED: the first thread can backtrack and stop.




Fig. 1 Using more than one thread for the exploration is useless.

A similar problem arise when performing on-the-fly model checking since (1) there is only one initial state and (2) all states are generated during the exploration. Thus a thread cannot be launched from a particular state. Moreover, the system's topology is only known after the exploration: we need a technique that works for any kind of topology.

The idea developed in this paper is the automatic generation of state x using genetic algorithms. The generation of *the perfect state* (the state x in the example) is a utopia. Nonetheless if we can generate a state relatively deep regarding to many DFS orders, we hope to avoid redundant work between threads, and thus maximize the information shared between threads. In practice we may generate states that do not belong to the state space, but Section 6 shows that more than 84% of generated states belongs to it.

3 Generation of Artificial Initial State

Genetic algorithms. For many applications the computation of an optimal solution is impossible since the set of all possible solutions is too large to be explored. To address this problem, Holland [12] proposed a new kind of algorithms (now called genetic algorithms) that are inspired by the process of natural selection. These algorithms are often considered as optimizer and used to generate high-quality solutions to search problems. Basically, genetic algorithms start by a population of candidate solutions and improve it using bio-inspired operators:

- *Crossover*: selects multiple elements in the population (the parents) and produces a child solution from them.
- *Mutation*: selects one element in the population and alters it slightly.

a	b
00101010	00110011

Fig. 2 Chromosome representation.

	Process 1	
	a	b
<i>parent1</i>	00000000	00000000
<i>parent2</i>	11111111	11111111
Crossover(S)	00000000	11111111

Fig. 3 Possible Crossover.

Applying and combining these operators produces a new generation that can be evaluated using a *fitness* function. This fitness function allows to select the best elements (w.r.t the considered problem) of this new population. These best elements constitute a new population on which mutation and crossover operations can be re-applied. This process is repeated until some satisfying solution is found (or until a maximal number of generations has been reached).

Genetic algorithms rely on a representation of solutions that is chromosome-like. In the definition of a transition system we observe that every state can be seen as a tuple of integer variables using the γ function. Each variable can be considered as a gene and the set of variables can be considered as a chromosome composed of 0 and 1. For instance, if a state is composed of two variables $a = 42$ and $b = 51$ the resulting chromosome (considering 8 bits integers) would be the one described Figure 2.

Crossover. Concurrent reactive systems are generally composed of a set of N_p processes and a set of shared variables (or channels). Given a transition system $T = \langle Q, \iota, \delta, V, \gamma \rangle$ we can define $E : V \rightarrow [0, N_p]$, such that if v is a shared variable, $E(v)$ returns 0 and otherwise $E(v)$ returns the identifier of the process where the variable v is defined.

Algorithm 2 defines the crossover operation we use. This algorithm takes a parameter S which represents the population to use for generating a new state. Line 2 instantiates a new state s that will hold the result of the crossover operation. Lines 3 to 5 set up the values of the shared variables of s : for each shared variable v , an element of S is randomly selected to be the parent. Then, at line 5, one can observe that $\gamma(s)[v]$ (the value of v in s) is set according to $\gamma(\text{parent})[v]$ (the value of v in the *parent*). Lines 6 to 9 perform a similar operation on all the remaining variables.

These variables are treated by batch, i.e., all the variables that belong to a same process are filled using only one parent (line 7). This choice implies that in our **Crossover** algorithm the local variable of a process cannot have two different parents: this particular

	Process 1	
	a	b
s	00000100	00001000
Mutation(s)	0000010 1	00001000

Fig. 4 Possible Mutation.

processing helps to exploit the concurrency of underlying system. A possible result of this algorithm is represented Figure 3 (with 8 bits integer variables, only one process, no shared variables, $S = \{parent1, parent2\}$ and $child$ the state computed by $Crossover(S)$).

Algorithm 2: Crossover.

```

1 Procedure Crossover( $S \subseteq Q$ )
2    $s \leftarrow newState()$ 
3   for  $v \in V$  s.t.  $E(v) = 0$  do
4      $parent \leftarrow pick\ random\ one\ of\ S$ 
5      $\gamma(s)[v] \leftarrow \gamma(parent)[v]$ 
6   for  $i \in [0, N_p]$  do
7      $parent \leftarrow pick\ random\ one\ of\ S$ 
8     for  $v \in V$  s.t.  $E(v) = i$  do
9        $\gamma(s)[v] \leftarrow \gamma(parent)[v]$ 
10  return  $s$ 

```

Algorithm 3: Mutation.

```

1 Procedure Mutation( $s \in Q$ )
2   for  $v \in V$  do
3      $r \leftarrow random(0..1)$ 
4     if  $r > THRESHOLD$  then
5        $\gamma(s)[v] = random\_flip\_one\_bit\_in(\gamma(s)[v])$ 
6        $\gamma(s)[v] = bound\_project(\gamma(s)[v])$ 

```

Mutations. The other bio-inspired operator simulates alterations that could happen while genes are combined over multiples generations. In genetic algorithms, these mutations are performed by switching the value of a bit inside of a gene. Here, all the variables of the system are considered as genes.

Algorithm 3 describes this mutation. For each variable in the state s (line 2), a random number is generated. A mutation is then performed only if this number is above a fixed threshold (line 4): this restriction limits the number of mutations that can occur in a chromosome. We can then select randomly a bit in the current variable v and flip it (line 5). Finally, line 6 exploits the information we may have about the system by restricting the mutated variable to its bounds.

Indeed, even if all variables are considered as integer variables there are many cases where the bounds are known a priori: for instance Boolean, enumeration

types, characters, and so on are represented as integers but the set of value they can take is relatively small regarding the possible values of an integer. A possible result of this algorithm is represented Figure 4 (with 8 bits integer variables and only two character variables, i.e., that have values between $[0..255]$).

Fitness. As mentioned earlier, every new population must be restricted to the only elements that help to obtain a better solution. Here we want to generate states that are (1) reachable and (2) deep with respect to many DFS orders. These criteria help the swarming technique by exploring parts of the state space before another thread (starting from the real initial state) reaches them.

We face here a problem that is: for a given state it is hard to decide whether it is a *good* candidate without exploring all reachable states. For checking *deadlocks* (i.e., states without successors) Godefroid and Khurshid [10] proposed a fitness function that will only retains state with few transitions enabled⁴.

Since we have different objectives a new fitness function must be defined. In order to maximize the chances to generate a reachable state, we compute the average outgoing transitions (T_{avg}) of all the states that belong to the initial population. Then the fitness function uses this value as a threshold to detect *good* states. Many fitness function can be considered:

- **equality:** the number of successors of a *good* state is exactly equal to T_{avg} . The motivation for this fitness function is that if there are $N > 1$ independent processes that are deterministic then at every time, any process can progress. In this strategy, we consider that a good state has exactly N (equal to T_{avg}) outgoing transitions.
- **lessthan:** the number of successors of a *good* state is less than T_{avg} . The motivation for this fitness function is that if there are $N > 1$ independent deterministic processes that communicate then at any time each process can progress or two processes can be synchronized. This latter case will reduce the number of outgoing transitions
- **greaterthan:** the number of successors of a *good* state is greater than T_{avg} . The motivation for this fitness function is that if there are $N > 1$ independent and non-deterministic processes then at any time each processes can perform the same amount of actions or more.

⁴ Godefroid and Khurshid [10] do not generate states but finite paths and their fitness function analyzes the whole paths to keep only those with few enabled transitions.

Algorithm 4: The generation of new states.

```

1 Procedure Generate( $S \subseteq Q$ )
2   for  $i \leftarrow 0$  to NB_GENERATION do
3      $S' \leftarrow \emptyset$ 
4     for  $j \leftarrow 0$  to POP_SIZE do
5        $s \leftarrow$  Crossover( $S$ )
6       Mutation( $s$ )
7       if Fitness( $s$ ) then  $S' \leftarrow S' \cup \{s\}$ 
8      $S \leftarrow S'$ 
9   return  $S$ 

```

Generation of artificial state. Algorithm 4 presents the genetic algorithm used to generate artificial initial states using the previously defined functions.

The only parameter of this algorithm is the initial population S we want to mutate: S is obtained by performing a swarmed bounded DFS and keeping trace of all encountered states. From the initial population S , a new generation can be generated (lines 4 to 8). At any time the next generation is stored in S' (lines 7 and 3). The algorithm stops after NB_GENERATION generations (line 2). Note that this algorithm can report an empty set according to the fitness function used.

4 State-Space Exploration with Genetic Algorithm

This section explains how Algorithm 1 can be adapted to exploit the generation of artificial initial states mentioned in the previous section. Algorithm 5 describes this parallel state-space exploration using genetic algorithm. The basic idea is to have a *collaborative portfolio* approach in which threads will share information about CLOSED states. In this strategy, half of the available threads runs a the DFS algorithm presented Section 2, while the other threads perform genetic exploration. This exploration is achieved by three steps:

1. Perform swarmed bounded depth-first search exploration that stores into a set \mathcal{P} all encountered states (line 7). This exploration is *swarmed*, so that each thread has a different initial population \mathcal{P} . (Our *bounded*-DFS differs from the literature since it refers DFS that stops after visiting N states.)
2. Apply Algorithm 4 on \mathcal{P} to obtain a new population \mathcal{P}' of artificial initial states (line 8).
3. Apply the DFS algorithm for each element of \mathcal{P}' (lines 9 to 11). When the population \mathcal{P}' is empty, just restart the thread with the initial state ι (see line 12).

One can note (line 1) that the *color* enumeration has been augmented with OPEN_GP. This new status may seem useless for now but allows to distinguish states that have been discovered by the genetic algorithm from those discovered by the traditional algorithm. In this algorithm OPEN_GP acts and means exactly the same than OPEN but: (1) this status is useful for the sketch of termination proof below and, (2) the next section shows how we can exploit similar information.

Algorithm 5: Parallel DFS Exploration using Genetic Algorithm.

```

1 enum color = { OPEN, OPEN_GP, CLOSED }
2 visited: hashmap of ( $Q, color$ )
3 stop  $\leftarrow \perp$ 
4 Procedure ParDFS_GP( $\langle Q, \iota, \delta, V, \gamma \rangle : TS, n : Integer$ )
5   DFS( $\iota, 1, OPEN$ ) || ... || DFS( $\iota, \lfloor \frac{n}{2} \rfloor, OPEN$ ) || DFS_GP( $\iota, \lfloor \frac{n}{2} \rfloor + 1$ ) || ... || DFS_GP( $\iota, n$ )
6 Procedure DFS_GP( $\iota \in Q, tid : Integer$ )
7    $\mathcal{P} \leftarrow$  Bounded_DFS( $\iota, tid$ ) // Swarmed exploration
   using tid as a seed
8    $\mathcal{P}' \leftarrow$  Generate( $\mathcal{P}$ ) // Described Algorithm 4
9   while  $\mathcal{P}'$  not empty  $\wedge \neg stop$  do
10     $s \leftarrow$  pick one of  $\mathcal{P}'$ 
11    DFS( $s, tid, OPEN_GP$ )
12   if  $\neg stop$  then DFS( $\iota, tid, OPEN$ )

```

Termination. Until now we have avoided mentioning one problem: there is no reason that a generated state is a reachable state. Nonetheless even if the state is not reachable, some of its successors (direct or not) may be reachable. Since the number of unreachable states is generally much larger than the number of reachable states, we have to ensure that Algorithm 5 terminates as soon as all reachable states have been explored.

First of all let us consider only threads running the DFS algorithm. Since this algorithm has already been prove (see. [24] for more details), only the intuition is given here. When all the successors of an OPEN state have been visited, this state is tagged as CLOSED. Since all CLOSED states are ignored during the exploration, each thread will restrict parts of the reachable state space. At some point all the states will be CLOSED: even if a thread is still performing its DFS procedure, all the successors of its current state will be marked CLOSED. Thus the thread will be forced to backtrack and stop.

The problem we may have with using genetic algorithm is that all the threads performing the genetic algorithm may be running while all the other ones are idle since all the reachable states have already been

visited. In this case, a running thread can see only unreachable states, i.e. OPEN_GP, or CLOSED ones. To handle this problem, a Boolean *stop* is shared among all threads (line 2). When this Boolean is set to \top all threads stop regardless the exploration technique used (line 11, Algorithm 1). We observe line 9 that the use of other artificial states is also stopped, and no restart will be performed (line 12). This Boolean is set to \top only when all the successors of the real initial state have been explored (line 17, Algorithm 1). Thus, one can note that even if a thread using the genetic algorithm visits first all reachable states it will stop all the other threads.

5 Checking Temporal Properties

5.1 Checking Safety Properties

5.1.1 The deadlock detection algorithm

Safety properties cover a wide range of properties: *deadlock freedom* (there is no state without successors), *mutual exclusion* (two processes execute some critical section at the same time), *partial correction* (the execution terminates in a state that does not satisfies the postcondition while the precondition of the run was satisfied), *etc.* One interesting characteristic of safety properties is that they can be checked using a reachability analysis (as described Section 2). Nonetheless, our genetic reachability algorithm (Algorithm 5) cannot be directly used to check safety properties. Indeed, if a thread (using genetic programming) reports an error we do not know if this error actually belongs to the state space.

Algorithm 6 describes how to adapt Algorithm 5 to check safety properties. To simplify things we focus on checking deadlock freedom, but our approach can be generalized to any safety property. This algorithm⁵ relies on both Algorithms 1 and 5. The basic idea is still to launch half of the threads from the initial state ι and the remaining ones from some artificial initial state (line 10).

- For a thread performing reachability with genetic algorithm the differences are quite few. When a deadlock state is detected (line 36) we just tag this state as DEADLOCK_GP rather than CLOSED. This new status is used to mark all states leading to a deadlock state. Indeed since we do not know if the state is a reachable one we cannot report immediately that a deadlock has been found. Moreover we

Algorithm 6: Parallel Deadlock Detection Using Genetic Algorithm.

```

1  enum color =
2     { OPEN, OPEN_GP, CLOSED, DEADLOCK_GP }
3  visited: hashmap of (Q, color)
4  stop ← ⊥
5  deadlock ← ⊥
6  Procedure
   ParDeadlockGP(⟨Q,  $\iota$ ,  $\delta$ , V,  $\gamma$ ⟩ : TS, n : Integer)
7     DeadlockDFS( $\iota$ , 1, OPEN) || ... ||
8     DeadlockDFS( $\iota$ , ⌊ $\frac{n}{2}$ ⌋, OPEN) ||
9     DeadlockDFS_GP( $\iota$ , ⌊ $\frac{n}{2}$ ⌋ + 1) || ... ||
10    DeadlockDFS_GP( $\iota$ , n)

11 Procedure
   DeadlockDFS( $s \in Q$ , tid : Integer, status : color)
12  if  $s \notin visited$  then
13    | visited.add( $s$ , status)
14  else if visited[ $s$ ] = CLOSED then
15    | return
16  todo ← shuffle(post( $s$ ), tid)
17  while ( $\neg stop \wedge \neg todo.isempty()$ ) do
18    |  $s' \leftarrow todo.pick()$ 
19    if  $s'$  is in the current recursive DFS stack then
20      | continue
21    if ( $s' \notin visited \vee visited[s'] \neq CLOSED$ ) then
22      |  $b \leftarrow (s' \in visited \wedge visited[s'] = DEADLOCK\_GP$ 
23         |  $\wedge status = OPEN)$ 
24      if b then
25        | deadlock ←  $\top$ 
26        | stop ←  $\top$ 
27        | break
28      DeadlockDFS( $s'$ , tid, status)
29      if
30         | visited[ $s'$ ] = DEADLOCK_GP  $\wedge$  status = OPEN_GP
31         | then
32           | visited[ $s$ ] ← DEADLOCK_GP
33           | return
34      if post( $s$ ) =  $\emptyset \wedge status = OPEN$  then
35        | deadlock ←  $\top$ 
36        | stop ←  $\top$ 
37        | return
38      if post( $s$ ) =  $\emptyset \wedge status = OPEN\_GP$  then
39        | visited[ $s$ ] ← DEADLOCK_GP
40      else
41        |  $v[s] \leftarrow CLOSED$ 
42        if ( $s = \iota$ ) then
43          | stop ←  $\top$ 

42 Procedure DeadlockDFS_GP( $\iota \in Q$ , tid : Integer)
43  // Also check deadlock during this DFS
44   $\mathcal{P} \leftarrow Bounded\_DFS(\iota, tid)$ 
45   $\mathcal{P}' \leftarrow Generate(\mathcal{P})$ 
46  while  $\mathcal{P}'$  not empty  $\wedge \neg stop$  do
47    |  $s \leftarrow$  pick one of  $\mathcal{P}'$ 
48    | DeadlockDFS( $s$ , tid, OPEN_GP)
49  if  $\neg stop$  then
50    | DeadlockDFS( $\iota$ , tid, OPEN)

```

⁵ Main differences have been highlighted to help the reader.

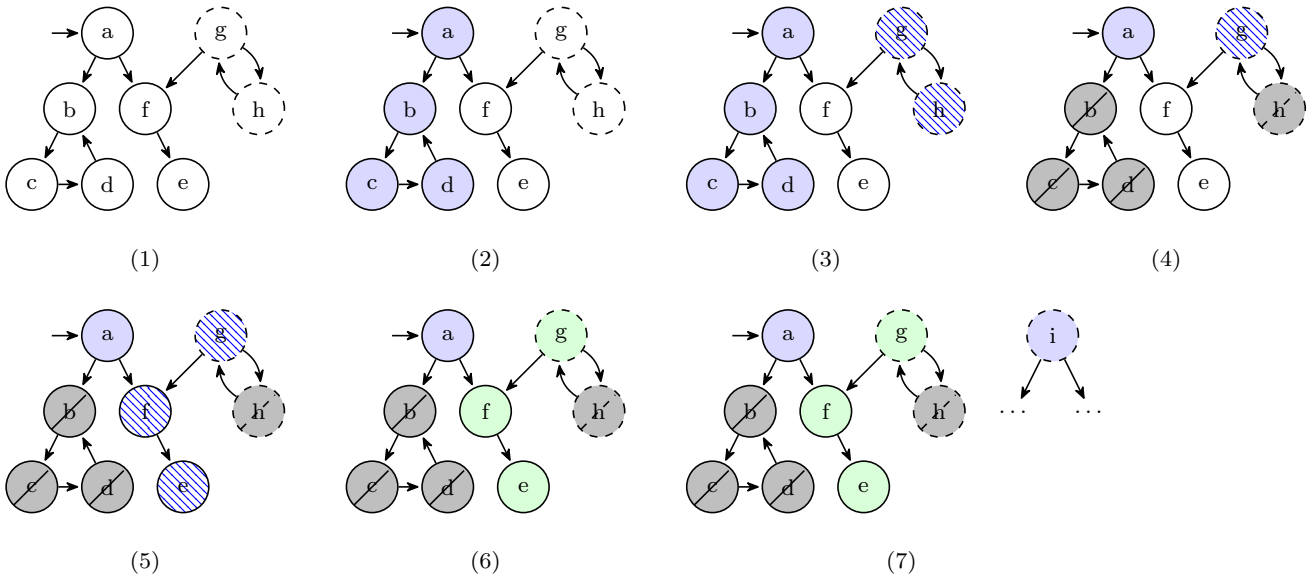


Fig. 5 Plain states represents reachable states while dashed ones represents non reachable ones. Two threads explore this example, one from a one from g (then from i). Blue states represent OPEN states, and blue cross hatched states represent OPEN_GP ones. Forbidden signed states represent CLOSED states and green ones the DEADLOCK_GP ones.

cannot mark this state CLOSED otherwise a counterexample could be lost. This new status helps to solve the problem: when such a state is detected to be reachable, a deadlock is immediately reported. The other modifications are lines 29 and 31: when backtracking, if a deadlock has been found no more states will be explored.

- For a thread performing reachability without genetic algorithm the differences are also quite few. Lines 24 to 27 only check if the next state to process has been marked DEADLOCK_GP. In this case this state is a reachable one and it leads to a deadlock state. We can then report that a deadlock has been found and stop all the other threads. A deadlock can also be reported directly (line 32), if the current state is a deadlock.

5.1.2 Detailed example.

Algorithm 6 is depicted step by step in Figure 5. For this example, we consider that lines 44 and 46 (computing the artificial initial states) have already been realized. Let us consider two threads: t_1 performing a *classical approach* and t_2 performing a *genetic programming* approach, i.e. $n = 2$, t_1 has been launched line 7, and t_2 line 10. Step (1) represents the state space (a, b, c, d, e, f) as well as some states that does not belong to it (g, h).

Step (2) represents t_1 starting from the initial state and exploring states a, b, c , and d . Step (3) represents t_2 starting from the artificial initial g state and exploring state h . In step (4), t_2 detects that all its successors

are on its recursive stack. As a consequence, the state h will be marked CLOSED, line 39 of Algorithm 6. The thread t_2 can then backtrack state h and continue to explore the next successors of g , i.e. f .

Step (5) represents two concurrent action. First, t_2 has explored states f and e . Second, t_1 has detected that all successors of d are on its recursive stack and then d has been marked as CLOSED. Then t_2 discovered that all the successors of c have been explored, so this state can also be marked CLOSED line 39. The same operations are then applied to b which is then marked CLOSED. t_1 can then backtracks b and explores the remaining successors of a .

During step (6), thread t_2 discovers that state e is effectively a deadlock (line 37). State e is then tagged DEADLOCK_GP. When this state is backtracked (after the recursive call line 28), state f detects that its only successors can reach a deadlock states. As a consequence, state f will be immediately marked DEADLOCK_GP (line 30), and backtracked (line 31). For the same reasons, states g will also be tagged DEADLOCK_GP, and backtracked. Notice that lines 40–41 prevents stopping all the threads when backtracking state g .

From now, as soon as t_1 will discover state f it can report a counterexample, i.e. a deadlock has been detected. Indeed, lines 21–27 of Algorithm 6 detects such a situation. This situation can be described as follow: ”a *classical thread* detects a state that can lead to a deadlock **but** discovered by a *gp thread*”. In this case, we can claim that there exist a path from the

initial state to a deadlock. One should note that t_1 reports a deadlock without seeing f and e . When the deadlock is detected, all the other threads are stopped.

Finally step (7) depicts lines 46–48 of the algorithm. Thread t_2 finished its exploration from state g and picked another artificial state (here state i).

Discussion. The observer reader may notice three relevant informations:

1. Here, the *gp thread* start two explorations, one from g and one from i . Both of these states does not belong to the state space. There is no obligation for these state to be outside of the state space. If they belong to the state space, the algorithm works perfectly the same, without this information.
2. Suppose that in step (3), the algorithm choose state f rather than state h . In this case, a deadlock will be found, and all states are backtracked. Doing that will prevent the exploration of state h . We opted for this strategy in order to propagate as soon as possible the information about deadlock detection. Nonetheless, our algorithm is easily adaptable to force the exploration of remaining successors.
3. Our algorithm does not exploit the fact that two *gp threads* cooperate. Indeed, a first thread can detect a deadlock and backtracks. When the second *gp thread* discovers a DEADLOCK_GP state it can immediately backtracks while our current algorithm force the exploration until the deadlock is re-discovered.

5.1.3 Proof of the algorithm

This subsection details the proof that Algorithm 6 will report a deadlock if and only if there exists a reachable state that has no successors. To prove this algorithm, two theorems must be verified:

Theorem 1. For all systems S , the algorithm terminates.

Theorem 2. A thread reports a deadlock iff $\exists s \in Q$, $post(s) = \emptyset$.

To simplify this proof, we denote by *classical thread* a thread that does not perform genetic algorithm while the other threads are called *gp threads*. The following invariants hold for all lines of Algorithm 6:

Invariant 1. If *stop* is \top then no new state will be discovered.

Proof. *New states are computed line 16 but only discovered one-by-one line 18. Let us suppose that*

some thread set the stop variable to \top (lines 26, 34, or 40): this thread will quit the while loop line 17. Exiting the DEADLOCKDFS function will also exit the loop line 45 and exit this thread. For the other threads two situations may occur. First, the threads are backtracking from the call line 28: the next iteration will not be executed, and the threads will exit without discovering new states. If the threads are executing lines 18–27, then the call line 28 will be performed but the check line 17 will avoid new states to be discovered.

Invariant 2. A deadlock state can only be OPEN, OPEN_GP or DEADLOCK_GP.

Proof. *From line 7–10, 47 and 49, the only status that can be used line 12–13 are OPEN and OPEN_GP. If a state has no successor, the condition if the loop (line 17) will not be satisfied and the thread jumps line 32. If the thread is using the classical approach, lines 32–34 are executed and a deadlock is reported (stopping the other threads). Otherwise, the state is only marked DEADLOCK_GP line 37.*

Invariant 3. No direct successor of a CLOSED state is a deadlock state.

Proof. *A state s is marked as CLOSED line 39 when all its successors have been visited lines 17 – 31. Lines 19 and 21 ensure that a recursive call is performed only on states that are (1) OPEN or not in the DFS stack. All the other direct successors are then explored and backtracked, and then marked CLOSED before s is marked CLOSED.*

Invariant 4. A state is marked CLOSED iff all its successors that are not on the thread's recursive stack are CLOSED.

Proof. *From Invariant 4, we know that all the direct successors of a state are either on the DFS stack or CLOSED. Since states are marked CLOSED in the DFS postorder line 39, all its successors that are not on the recursive stack are backtracked and then marked CLOSED.*

Invariant 5. Only *gp threads* can tag a state DEADLOCK_GP.

Proof. *Trivial. From the algorithm, the only places where the status is changed to DEADLOCK_GP are lines 30 and 37. For both, the previous line checks wheter the status is OPEN_GP. From line 9,10 and 48 only gp threads can have this status. One should note that line 50, the gp thread becomes a classical thread.*

Invariant 6. A state is DEADLOCK_GP iff it is a deadlock state or if one of its successors (direct or not) is a deadlock state.

Proof. If a state is trivially a deadlock state, line 37 will mark it `DEADLOCK_GP` and return (line 35). Consequently, all its predecessors will be marked `DEADLOCK_GP` during the backtrack (line 30). Thus, a state can only be tagged `DEADLOCK_GP` iff one of its successors is a real deadlock.

Invariant 7. Only classical thread can report that a deadlock has been found.

Proof. A deadlock can be reported lines 25–26 or lines 33–34. Lines 33–34 can only be executed by a classical thread due to the condition line 32 (and the note of invariant 5). Lines 25–26 can also report that state s is a deadlock but the condition ($\text{status} = \text{OPEN}$) ensures that only a classical thread will report it. In this case a *gp* thread has discovered that some (possibly indirect) successor of s is a deadlock, without knowing that s is reachable from the initial state. The lines 25–26 detect that this state is reachable and can then report the deadlock.

Invariant 8. If a state is reachable then all its direct successors are reachable.

Proof. By construction, ι is the initial state then reachable. All threads starting from ι lines 7,8 and 50 will then start from a reachable state. Applying the transition relation line 16 will then only produce reachable states.

Proof of Theorem 1. From invariant 8, and since the system has a finite number of states, at least one thread will perform the exploration from the initial state ι . If no deadlock is found, the thread will backtrack and finally reach lines 40–41. The stop boolean will then be set to \top . From invariant 1. the algorithm stops. Reaching `CLOSED` states (line 14) will only prune the exploration and then have no impact on the terminaison for classical threads. *Gp* threads may only explore states that are not part of the state space. Nonetheless, invariant 1 ensures the terminaison for these threads when some thread will mark stop as \top . If a deadlock is discovered by a classical thread, invariant 1 and 4 will ensures the terminaison of all threads. If a *gp* threads detects it, invariant 7 combined to invariant 8 will force all threads to stop using lines 26 and 34.

Proof of Theorem 2. From invariant 2 we know that a deadlock state can never be marked `CLOSED` because (1) if discovered by a classical thread a deadlock is immediately reported, and (2) because otherwise this information must be propagated. Invariant 3, 5 and 6 ensure that the information is correctly propagated, while invariant 7 ensure that no *gp* thread can report

that a deadlock has been found. The other direction of theorem 2 is quite evident. If a deadlock exist, then it is a reachable state. Since all reachable states are explored by classical thread, the report will be done (see invariant 7)

5.1.4 Reporting counterexample

In Algorithm 6, a classical thread can report the existence of a deadlock but cannot report the counterexample forming it. Indeed, the path from the initial state to the deadlock may be composed of several parts, computed by one *classical thread* and multiple *gp threads*.

Reporting the counterexample can be done as following. First of all the recursive call stack forms the prefix, starting from the initial state ι to some state α . Note that we know that this prefix starts from ι since only *classical threads* can report deadlocks.

Then, the thread must compute the path from α to one deadlock state β . To do so, the thread will compute the successors of α and choose one of them which is tagged `DEADLOCK_GP`. This operation is then repeated from the chosen successor. Indeed, after a *gp thread* has detected a deadlock, all the states on its DFS stack are tagged `DEADLOCK_GP`. Following a path of `DEADLOCK_GP` states will necessarily results in discovering a deadlock.

Combining the prefix and the path of `DEADLOCK_GP` states will build the whole counterexample.

5.2 Checking Liveness Properties

Until now, we only focused on the verification of safety properties. The verification of complex temporal properties involves the exploration of an automaton which is the result of the synchronous product between the state space of the system and the property automaton.

In this settings, a state is composed of two parts: the system state and the property state. Thanks to the previous sections, we know how to build artificial states for the system part.

Generating artificial initial states for the property may be irrelevant. Indeed, these automata may have a huge impact (in term of states) on the synchronized product. Moreover, we know that only states synchronized with some real state of the property automaton will report a counterexample. The generation of artificial states for the property automaton may not be relevant.

Nonetheless, genetic algorithms presented so far can then be applied by considering that the property

state is a variable just like the other system’s variables. For a three states property automaton, we can consider that the actual state in the property automaton depends on a variable that can have three values: 1, 2 and 3⁶.

The adaptation of the artificial state generation is then straightforward: the system part is generated as previously while the property part is generated as described earlier.

One should note that the generation of artificial initial state for the synchronized product is not sufficient for adapting Algorithm 6 for checking liveness properties. Indeed, checking liveness properties involves the use of an emptiness check in the automata theoretic approach for explicit LTL model-checking. An emptiness check is an algorithm looking for *accepting cycles* in the synchronized product, i.e. cycles that contains products states synchronized with some designated state of the property automaton.

Traditionally, two kind of emptiness checks are used in explicit model-checking⁷:

- NDFS-based [8]: two nested dfs are used to detect accepting cycles. A first one looks for the accepting state, while the second one looks for a cycle around it.
- SCC-based [24]: one dfs is used to compute the *Strongly Connected Components* (SCC) of the synchronized product. As soon as a an SCC containing an accepting state is discovered a counterexample can be reported.

The adaptation of Algorithm 6 into these emptiness check can be done as following.

- For NDFS-based algorithms, when a *gp thread* detects an accepting cycle, all the states forming it are tagged with an `ACCEPTING_CYCLE` status. When a *classical thread* detects such a state, a counterexample is raised.
- For SCC-based algorithms, when a *gp thread* detects an accepting SCC, all the states forming it are tagged with an `ACCEPTING_SCC` status. When a *classical thread* detects such an SCC, a counterexample is raised. One should note that tagging all the states of an SCC can be done in quasi-constant time using a union-find data structure (see Anderson and Woll [2] for a lock-free implementation of this structure).

⁶ Notice that mutation can be done ensuring that this variable will not be less than 1 and not be greater than 3

⁷ Here, we only describe our approach on Büchi automata, but the adaptation for generalized Büchi automata is straightforward

Notice that in both situations, as soon as a *classical thread* detects a counterexample, this latter one can be immediately reported and all the other threads stopped.

6 Evaluation

Benchmark Description. To evaluate the performance of our algorithms, we selected 38 models from the BEEM benchmark [18] that cover all types of models described by the classification of Pelánek [19]. All the models were selected such that Algorithm 1 with one thread would take at most 40 minutes on Intel(R) Xeon(R) @ 2.00GHz with 250GB of RAM. This six-core machine is also used for the following parallel experiments⁸. All the approaches proposed here have been implemented in Spot [7]. For a given model the corresponding system is generated on-the-fly using DiVinE 2.4 patched by the LTSmin team⁹.

Reachability. To evaluate the performance of the algorithm presented Section 4 we conducted 9158 experiments, each taking 30 seconds on the average. Table 2 reports selected results to show the impact of the fitness function and the threshold over the performance of Algorithm 5 with 12 threads (the maximum we can test). For each variation, we provide *nb* the number of models computed within time and memory constraints, and *Time* the cumulated walltime for this configuration (to run the whole benchmark). For a fair-comparison, we excluded from *Time* models that cannot be processed. Table 2 also reports state-of-the-art and **random** (used to evaluate the accuracy of genetic algorithms by generating random states as seed state). This latter technique is irrelevant since it is five time slower than state-of-the-art and only process 32 models over 38.

If we now focus on genetics algorithms, we observe that the threshold highly impacts the results regardless the fitness function used: the more the threshold grows, the more models are processed within time and memory constraints.

The table also reports the best threshold¹⁰ for all fitness function, i.e. 0.999. It appears that **greaterthan**

⁸ For a description of our setup, including selected models, detailed results and code, see <http://www.lrde.epita.fr/~renault/benchs/VECOS-2018/results.html>

⁹ See <http://fmt.cs.utwente.nl/tools/ltsmin/#divine> for more details. Also note that we added some patches (available in the webpage) to manage out-of-bound detection.

¹⁰ We evaluate other thresholds like 0.9999 or 0.99999 but it appears that augmenting the threshold does not increase performance, see the webpage for more details.

	THRESHOLD							
	0.7		0.8		0.9		0.999	
	<i>nb</i>	<i>Time (ms)</i>	<i>nb</i>	<i>Time (ms)</i>	<i>nb</i>	<i>Time (ms)</i>	<i>nb</i>	<i>Time (ms)</i>
greaterthan	35	1 041 015	35	970 248	35	1 000 184	37	900 468
equality	35	3 217 183	35	965 259	35	934 947	38	907 148
lessthan	35	972 038	35	951 767	35	928 978	38	904 776
lessstrict	35	970 668	35	983 225	35	935 319	38	894 131
	No threshold							
random	(trivial comparator to evaluate genetic algorithms)						32	5 079 869
Algorithm 1	(state-of-the-art with 12 threads)						38	978 711

Table 2 Impact of the threshold and the fitness function on Algorithm 5 with 12 threads (`NB_GENERATION=3`, `INIT=1000`, `POP_SIZE=50`). The time is expressed in millisecond and is the cumulated time taken to compute the whole benchmark (38 models); *nb* is the number of instances resolved with time and memory limits.

only processed 37 models: this fitness function does not seem to be a good fitness function since (1) it tends to explore useless parts of the state-space and (2) the variations of the threshold highly impacts the performance of the algorithm. All the other fitness function provide similar results for a threshold fixed at 0.999. Nonetheless we do not recommend **equality** since a simple variation of the threshold (0.7) could lead to extremely poor results. Our preference goes to **lessthan** and **lessstrict** since they seem to be less sensitive to threshold variation while achieving the benchmark 9% faster than state-of-the-art algorithm. Thus, while the speedup for 12 threads was 3.02 for state-of-the-art algorithm, our algorithm achieves a speedup of 3.31.

Note that the results reported Table 2 include the computation of the artificial initial states. On the overall benchmark, this computation take in average slightly less than 1 second per model (30 seconds for the whole benchmark). This computation has a negligible impact on the speedup of our algorithm.

We have also evaluated (not reported here, see webpage for more details) the impact of the size of the initial population and the size of each generation over the performance. It appears that augmenting (or decreasing) these two parameters deteriorate the performance. It is worth noting that the best value of all parameters are classical values regarding to state-of-the-art genetic algorithms. Finally, for each model (and **lessthan** as fitness), we compute a set of artificial initial states and run an exploration algorithm from each of these states. It appears that 84.6% of the 7 866 005 486 generated states are reachable states.

The chart presented Figure 6 evaluates the percentage of reachable states from a population of artificial initial states (computed with the best parameter inferred from Table 2). The results are presented model per model. For each model and for a given arti-

cial initial state we evaluated the percentage of visited states that are reachable from the real initial state. For a model, the boxplot displays this percentage from each artificial state in the population. Models are presented sorted according to their median value.

First of all, we can observe that all almost all models have at least one artificial state where all its successors are reachable from the initial state. Moreover one third of the models have more than a half of their artificial initial states with more than 50 percents of successors that are reachable. One can observe huge variations depending on models: for instance, almost all the states are reachable from any artificial state in *resistance.2* while there are few for *blocks.3*. A fine grained study of these model reveals that models with good results fall into two categories of the classification of Pelánek [19]: Mutex and Communication-Protocol. These models appears to be composed of large SCCs or long cycles. This chart suggests that the function used for the generation of artificial initial state may be crucial for our algorithms but may also be dependent of the kind of model targeted.

Safety properties. Now that we have detected the best values for the parameters of the genetic algorithm we can evaluate the performance of our deadlock detection algorithm. In order to evaluate the performance of our algorithm we conduct 418 experiments. The benchmark contains 21 models with deadlocks and 17 models without. Table 3 compares the relative performance of state-of-the-art algorithm and Algorithm 6. For this latter algorithm, we only report the two fitness functions that give the best performance for reachability. Indeed, since Algorithm 6 is based on Algorithm 5 we reuse the best parameters to obtain the best performance. Results for detecting deadlocks are quite disappointing since our algorithm is 15% to 30% slower. A closer look to these results show that deadlocks are detected quickly and Algorithm 6 has




Fig. 6 Distributions of reachable states from a population of artificial initial states.

degraded performance due to the computation of artificial initial states.

On the contrary we observe that our algorithm is 10% faster (regardless whether we use **lessthan** or **lesstrict**) than the classical algorithm when the system has no deadlock. One can note that this algorithm performs better than simple reachability algorithm. Indeed, even if the system has no deadlock: the algorithm can find non-reachable deadlock. In this case, the algorithm backtracks and the next generation is

processed. This early backtracking force the use of a new generation that will helps the exploration of the reachable states. To achieve this speedup, we observe an overhead of 13% for the memory consumption. The use of dedicated memory reduction techniques could help to reduce this footprint.

Variations of the numbers of gp threads. All the algorithms presented in this paper suppose that only half of the threads performs an exploration based

	Algorithm 1 (state-of-the-art)		Algorithm 6			
	Time (ms)	States	lessthan		lessstrict	
	Time (ms)	States	Time (ms)	States	Time (ms)	States
Deadlocks	2 888	$7.01E^6$	3 713	$5.87E^6$	3 414	$5.47E^6$
No deadlocks	516 152	$5.79E^8$	462 881	$6.73E^8$	468 683	$6.82E^8$

Table 3 Comparison of algorithms for deadlock detection. Each runs with 12 threads, and we report the variation of two different fitness functions: **lessstrict** and **lessthan**. Results presents the cumulated time and states visited for the whole benchmark.

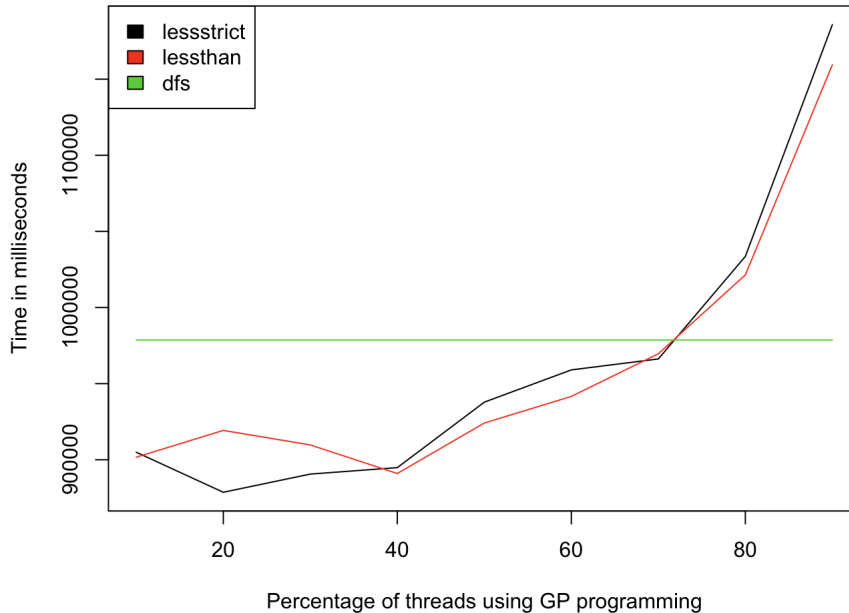


Fig. 7 Impact of the percentage of thread using genetic programming. Experiments were run with 12 threads. The *dfs* line represents state-of-the art while *lessthan* and *lessstrict* represents the fitness functions combined to the best parameters observed Table 2.

on one or more artificial states. This restriction was chosen based on the experiments Table 1 where half of the threads seems useless from the speedup point of view. Figure 7 describes possible variations on this approach. The main idea is to run Algorithm 5 with a variation on the percentage of threads using a genetic programming approach. This Figure displays three lines: *dfs* represents the state-of-the art while the other lines represent the fitness functions combined to the best parameters observed Table 2. It appears that the percentage of thread using a genetic approach have a strong impact on the results. Indeed, even with a few percentage (20%) of threads using genetic programming, we observe a 11% reduction of the total time to run the benchmark. Nonetheless we can observe that these performances are degraded while the percentage of *gp threads* augments. Indeed, *gp threads* may explore states that are not in the state space and using too much of these threads will not help *classical threads* to explore concurrently this state space. No-

tice that compared to results presented Table 2, the variation of the percentage of the *gp threads* helps to obtain around 2.5% more reduction of the total time for this benchmark.

Discussion. Few models in the benchmark have a linear topology, which can be considered as the perfect one for the algorithms presented in this paper. Nonetheless, we observe a global improvement of state-of-the-art algorithm. We believe that other fitness function (based on interpolation or estimation of distribution) could help to generate better states, i.e. deep with respect to many DFS orders.

7 Conclusion

We have presented some first and new parallel exploration algorithms that rely on genetic algorithms. We suggested to see variables of the model as genes and states as chromosomes. With this definition we

were able to build an algorithm that generates artificial initial states. To detect if such a state is relevant we proposed and evaluate various fitness functions. It appears that these seed states improve the swarming technique. This combination between swarming and genetic algorithms has never been proposed and the benchmark show encouraging results (10% faster than state-of-the-art). Since the performance of our algorithms highly relies on the generation of good artificial states we would like to see if other strategies could help to generate better states. We also observed that a small percentage of threads using genetic programming is sufficient to obtain a good speedup.

We also demonstrate the correctness of our algorithms and describe how our they can be adapted to report counterexamples.

This work mainly focused on checking safety properties even if we proposed an adaptation for liveness properties. A future work would be to evaluate the performance of our algorithm in this latter case. We also want to investigate the relation between artificial state generation and POR, since both rely on the analysis of processes variables. Finally, we strongly believe that this paper could serve as a basis for combining parametric model-checking with neural network.

References

1. P. E. Ammann, P. E. Black, and W. Majurski. Using model checking to generate tests from specifications. In ICFEM'98, pp. 46–54, december 1998.
2. R. J. Anderson and H. Woll. Wait-free parallel algorithms for the union-find problem. In Proc. 23rd ACM Symposium on Theory of Computing, pp. 370–380, 1994.
3. J. Barnat, L. Brim, and P. Ročkal. Scalable shared memory LTL model checking. *STTT*, 12(2):139–153, 2010.
4. V. Bloemen and J. van de Pol. Multi-core SCC-Based LTL Model Checking, pp. 18–33. Lecture Notes in Computer Science. Springer, 2016.
5. J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. Hwang. Symbolic model checking: 10^{20} states and beyond. In Proc. of the Fifth Annual IEEE Symposium on Logic in Computer Science, pp. 1–33, 1990. IEEE.
6. C. Courcoubetis, M. Y. Vardi, P. Wolper, and M. Yannakakis. Memory-efficient algorithm for the verification of temporal properties. In CAV'90, vol. 531 of LNCS, pp. 233–242. Springer, 1991.
7. A. Duret-Lutz, A. Lewkowicz, A. Fauchille, T. Michaud, E. Renault, and L. Xu. Spot 2.0 — a framework for LTL and ω -automata manipulation. In ATVA'16, vol. 9938 of LNCS, pp. 122–129. Springer, Oct. 2016.
8. S. Evangelista, A. Laarman, L. Petrucci, and J. van de Pol. Improved multi-core nested depth-first search. In ATVA'12, vol. 7561 of LNCS, pp. 269–283. Springer, 2012.
9. H. Garavel, R. Mateescu, and I. Smarandache. Parallel State Space Construction for Model-Checking. Technical Report RR-4341, INRIA, 2001.
10. P. Godefroid and S. Khurshid. Exploring Very Large State Spaces Using Genetic Algorithms, pp. 266–280. Springer, Berlin, Heidelberg, 2002.
11. P. Godefroid, G. J. Holzmann, and D. Pirottin. State space caching revisited. In CAV'92, vol. 663 of LNCS, pp. 178–191. Springer, 1992.
12. J. H. Holland. Genetic algorithms. *Scientific American*, 1992.
13. G. J. Holzmann. On limits and possibilities of automated protocol analysis. In PSTV'87, pp. 339–344. North-Holland, May 1987.
14. G. J. Holzmann and D. Bosnacki. The design of a multi-core extension of the SPIN model checker. *IEEE Transaction on Software Engineering*, 33(10):659–674, 2007.
15. G. J. Holzmann, R. Joshi, and A. Groce. Swarm verification techniques. *IEEE Transaction on Software Engineering*, 37(6):845–857, 2011.
16. G. Katz and D. A. Peled. Synthesis of parametric programs using genetic programming and model checking. In INFINITY'13, pp. 70–84, 2013.
17. A. Laarman, E. Pater, J. Pol, and H. Hansen. Guard-based partial-order reduction. *STTT*, pp. 1–22, 2014.
18. R. Pelánek. BEEM: benchmarks for explicit model checkers. In SPIN'07, vol. 4595 of LNCS, pp. 263–267. Springer, 2007.
19. R. Pelánek. Properties of state spaces and their applications. *International Journal on Software Tools for Technology Transfer (STTT)*, 10:443–454, 2008.
20. R. Pelánek, T. Hanzl, I. Černá, and L. Brim. Enhancing random walk state space exploration. In FMICS'05, pp. 98–105. ACM Press, 2005.
21. D. Peled. Combining partial order reductions with on-the-fly model-checking. In CAV'94, vol. 818 of LNCS, pp. 377–390. Springer, 1994.
22. J. H. Reif. Depth-first search is inherently sequential. *Information Processing Letters*, 20:229–234, 1985.
23. E. Renault. Improving parallel state-space exploration using genetic algorithms. In VECOS'18, vol. 11181 of LNCS, pp. 133–149, 2018. Springer.
24. E. Renault, A. Duret-Lutz, F. Kordon, and D. Poirineaud. Variations on parallel explicit model checking for generalized Büchi automata. *International Journal on Software Tools for Technology Transfer (STTT)*, pp. 1–21, Apr. 2016.
25. H. Sivaraj and G. Gopalakrishnan. Random walk based heuristic algorithms for distributed memory model checking. *Electronic Notes in Theoretical Computer Science*, 89(1):51 – 67, 2003.
26. A. Valmari. Stubborn sets for reduced state space generation. In ICATPN'91, vol. 618 of LNCS, pp. 491–515, 1991. Springer.