LEARNING TO GENERATE RELIABLE BROADCAST ALGORITHMS

Diogo Vaz, David R. Matos, Miguel L. Pardal, Miguel Correia
INESC-ID, Instituto Superior Técnico, Lisboa – Portugal
{diogo.vaz, david.r.matos, miguel.pardal. miguel.p.correia}@tecnico.ulisboa.pt

July 29, 2022

ABSTRACT

Modern distributed systems are supported by fault-tolerant algorithms, like *Reliable Broadcast* and *Consensus*, that assure the correct operation of the system even when some of the nodes of the system fail. However, the development of distributed algorithms is a manual and complex process, resulting in scientific papers that usually present a single algorithm or variations of existing ones. To automate the process of developing such algorithms, this work presents an intelligent agent that uses *Reinforcement Learning* to generate correct and efficient fault-tolerant distributed algorithms. We show that our approach is able to generate correct fault-tolerant *Reliable Broadcast* algorithms with the same performance of others available in the literature, in only $12,000$ learning episodes.

1 Introduction

Distributed systems are made of multiple components interconnected by communication networks. By leveraging distribution, such systems can provide high availability and scalability. Examples of modern and widely used distributed systems are cloud applications [3] and blockchains [50]. During normal operation, some of these components may fail, e.g., due to power loss, software bugs, or malicious attacks. These faults can compromise the normal functioning of the entire system. Therefore, it is necessary to provide fault tolerance properties so that the distributed system can maintain its normal execution, even in the presence of faults. For that purpose, it is important to design and implement *fault-tolerant distributed algorithms* [12, 42].

Fault-tolerant distributed algorithms have been widely studied and developed over the years [7–9, 14, 18, 19, 29], exploring different aspects such as the problem to study [14, 29], the failures mode [19, 26] or the system architecture [8, 18]. However, the process of studying and designing a fault-tolerant algorithm is a *manual* and *complex* endeavor [26]. This is especially true in the presence of malicious actors, where the algorithms are very subtle, complicated and slight changes often require a complete redesign of the algorithm. Moreover, current research on fault-tolerant algorithms do not focus only on *correctness*: *efficiency* is also a very important concern in these algorithms, which increases the complexity of the overall algorithm development process.

Typically, the journey to develop a fault-tolerant algorithm starts with assumptions about the environment or system model, e.g., if the system is synchronous or asynchronous and the failures that can occur. Researchers then have to define the problem they want to solve, e.g., Reliable Broadcast, shortened in this paper as RBcast. Next, researchers think about the strategy to design the distributed algorithm. In this stage, besides the difficulty of creating the algorithm, researchers may be biased by previous related papers and algorithms. After a trial-and-error process, the distributed algorithm is generated. This also involves a validation process to assess whether the generated algorithm achieves the goal and solves the problem correctly. This can be done by writing a proof and/or by doing verification using a model checker or a theorem prover.

In this work, our aim is to automate the process of generating fault-tolerant distributed algorithms by proposing an *intelligent agent* [48] capable of generating and validating algorithms for a specific distributed problem.[1] More precisely, we aim to create an agent that can generate correct and efficient algorithms based on the *inputs* given by researchers, i.e., assumptions about the environment, failures, and characteristics of the algorithm. Our solution aims to be *flexible* so that researchers can change the specification with the objective of studying different problems that, consequently, can lead to different algorithms. This is a novel research area, so in this paper we start with a single problem: RBcast, but our work aims to be applicable to other distributed problems.

Our approach follows a process with two phases: *generation*, for obtaining candidate algorithms, and *validation*, for assessing their correctness. The *generation* process is related to the problem of code/program generation but with significant differences, as distributed algorithms are executed in parallel in several nodes and are subject to faults, whereas local programs are not. Research on automatic code generation explored static techniques such as design patterns [11], UML [46], and reverse engineering [49] but lately, supervised machine learning techniques have started to be used [1, 4, 57] to improve the process. In this work, we use an *unsupervised machine learning* approach called Reinforcement Learning [32, 54, 56]. To the best of our knowledge, this work is the first to present such an approach on the automatic code generation field.

In terms of *validation*, the search for a validation process for distributed algorithms has been pursued for years [23, 25, 38, 51]. We identified four possible languages and frameworks to be used: the TLA+ language and tools [38], the Spin framework with the PROMELA language [31], the ByMC framework [35], and IC3PO [24]. We opted to use Spin/PROMELA, as explained later.

The main contributions of the paper are: (1) a new approach for generating correct and efficient distributed fault-tolerant RBcast algorithms using machine learning, instead of manual development by human beings, (2) an intelligent agent to generate such distributed algorithms and (3) an experimental evaluation of the approach and the agent, showing a correct generation of RBcast algorithms.

2 Reliable Broadcast

This section presents RBcast, the distributed problem for which we want to find algorithms to solve.

2.1 System Model

The system model considered for the RBcast algorithms we want to generate is inspired on the modular approach to fault-tolerant problems by Hadzilacos and Toueg [26]. The system is composed of a static group of N processes, i.e., there are no joins or leaves during execution. We assume a *fully connected point-to-point* network, i.e., that all processes are connected with each other through *links* and communicate by passing messages. We also assume that the system is *asynchronous*, i.e., that the communication delays are neither upper bounded nor respect a global stabilization time.

A *process* is the actor of the distributed system, that executes a set of specific ordered *actions*, designated an *algorithm*. All processes of the system execute the same algorithm.

The communication links allow processes to exchange *messages*. The links have the task of transporting the message from the output buffer of the sender to the input buffer of the receiver. We assume that the links are *reliable*, *authenticated*, and *provide integrity on the messages*, meaning that there are no corrupted, lost or duplicated messages. However, messages may arrive out of order.

Processes use *messages* to share data between them. Typically, a message contains data such as the message content (used by the logic of the system) and an identifier that can contain *protocol type*, *sender*, and *sequence number*.

The processes of the system can be correct or faulty. We consider *three failure modes*: No-Failure, Crash-Failure, and Byzantine-Failure. In the simplest, *No-Failure*, we assume there are no failures. In the *Crash-Failure mode* [6, 21], processes may stop operating and never recover. Assuming a system with $N \in \mathbb{N}$ processes, the maximum number of faulty processes $F \in \mathbb{N}$ due to a crash failure that can be tolerated in the system is $F = \lfloor (N - 1)/2 \rfloor$ [10]. In the *Byzantine-Failure mode* [39] faulty processes may have arbitrary behavior, e.g., they may execute other actions not defined by the algorithm or even not execute any action at all. Unlike a crash failure, when a process suffers a Byzantine failure (or, "is Byzantine"), it can continue to work. Assuming a system with N processes, the maximum number of faulty processes F due to a Byzantine failure that can be tolerated in the system is $F = \lfloor (N-1)/3 \rfloor$ [10,18].

---

[1]We take the term *intelligent agent* from the Artificial Intelligence area, where Reinforcement Learning appeared. If the term "intelligent" makes sense is, of course, debatable.

2.2 Problem Definition

A RBcast algorithm ensures, essentially, that every message broadcast by a *correct* process (an *RB-Broadcast* event) is eventually delivered by all *correct* processes. The protocol is defined by the following properties [13, 26]:

- *RB-Agreement*: if a correct process delivers a message *m*, then all correct processes will eventually deliver the same message *m*;

- *RB-Validity*: if a correct process broadcasts a message *m*, then it will eventually deliver that message *m*;

- *RB-Integrity*: for any message *m*, every correct process delivers *m* at most once and only if *m* was previously broadcast by some correct process.

The term *correct* refers to a process that follows the algorithm. Otherwise, we call it incorrect or faulty.

2.3 Algorithms

RBcast algorithms can differ. In this section, we present the definitions of the algorithm adopted in this paper, namely, the structure, messages, types, conditions and efficiency.

2.3.1 Structure

RBcast has been much studied over the years [7–9, 18, 29, 30]. The papers present RBcast algorithms with different structures. For example, Bracha and Toueg presented algorithms organized in execution steps [9, 10], whereas more recent work favors a structure in terms of event handling routines [29, 30]. We follow the latter, event-oriented, structure.

We assume that the structure of the algorithm is composed of two events: the *RB-Broadcast* event, triggered only once by the process that starts the execution of the algorithm; and a *receive* event, triggered every time a process receives a message. Each event can contain a set of *actions*, i.e., instructions that are executed when the event is triggered. The execution of the algorithm ends when there are no more messages to be received.

2.3.2 Messages and Types

As previously discussed in Section 2.1, messages typically contain multiple elements, e.g., content, sender, and protocol type. However, to solve the RBcast problem, it is necessary to add a new element, called *type*, to differentiate between the different communication steps of the algorithm. In this work, we follow the definition of previous RBcast algorithms [9, 30], by presenting the message in the format <*t,m*>, where *t* symbolizes the type of the message, and *m* the message itself. Previous works define the types as words like *echo* or *init* [9, 29], but in this work, we designate types *type0*, *type1*, *type2*, etc. The reason for this option is that the types are generated automatically, not by a human. But afterwards, we can translate the type *type0* into *init* and the type *type1* into *echo*.

2.3.3 Conditions

Conditions are statements used to evaluate when a specific action can be executed and are associated with the *if* clause. On RBcast, a condition is defined by two properties: the *message* <*t,m*> and the *threshold*, which is the number of messages needed to satisfy the condition. For example, the condition *if received (<type0,m>) from $F + 1$ distinct processes* means that the action can only be executed if the process has received a message *m* of *type0* from at least $F + 1$ processes (the threshold is $F + 1$).

2.3.4 Efficiency

To analyze the efficiency of an algorithm, we adopt a model, based on previous works [12, 52], where the efficiency is related to three properties: (1) the number of messages sent by the algorithm; (2) the number of communication steps (which in the case of RBcast is the number of types, as there are no loops) and; (3) the number of messages that have to be received for the algorithm to stop. All these metrics indirectly express the cost of computational power, storage and network to execute the algorithm. As the algorithm sends more messages or contains more communication steps, then the process will need more storage for messages, spend more network resources, and take more time to execute the algorithm.
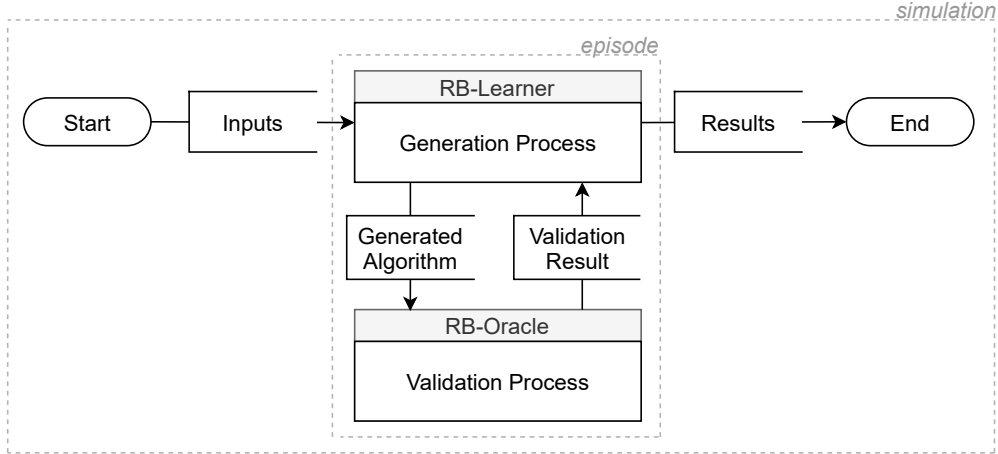
Figure 1: Process/dataflow of generating one algorithm. One simulation involves many episodes.

Table 1: Simulation inputs.

Generation process inputs	Possible Values	Validation process inputs	Possible Values
Number of simulations	{1,2,3,...}	Failure modes to model	{No-Failure,Crash-Failure, Byzantine-Failure}
Number of episodes	{1,2,3,...}	Number N of nodes to model	{3,4,5,...}
Rewards	{...,-2,-1,0}	Fault tolerance ratios	{0,(N-1)/2, (N-1)/3}
			for No-/Crash-/Byzantine-Failure, respectively
Heuristics to be applied	(cf. Sec. 4.5)	Properties to be validated	{RB-Validity, RB-Agreement, RB-Integrity}
		Event handlers where faults can occur	{RB-Broadcast, receive}

3 Learning the RBcast Algorithm

Machine learning techniques can be divided into two broad classes: *supervised machine learning* – to learn from labeled data – and *unsupervised machine learning* – to learn without labeled data. In this task of learning how to solve a distributed problem, we decided to emulate the *trial and error* strategy used by a human researcher to solve this problem. Therefore, we decided to use an unsupervised machine learning technique called *Reinforcement Learning* [32, 54, 56].

Reinforcement Learning is based on the idea of an *agent* choosing *actions* in specific *states*. The states are the observations that the agent receives from the *environment* where it acts. The way in which the agent acts is defined by the *policy*, in this case, a map from perceived states to actions to be taken in those states. Then, by choosing an action in a state, the agent will receive a *reward*, reflecting its choice. With time, the agent will start learning the *value* of each state, i.e., the total amount of reward the agent can expect to accumulate over the future, starting from that state. While rewards are *short-term* indicators, the values reflect the *long-term* desirability of that state, taking into account the states that are likely to follow and the rewards associated with them. We consider a *model-free* Reinforcement Learning problem, meaning that the agent does not create a representation of the behavior of the environment, unlike model-based approaches. A model-based agent would needed to have a model of the dynamics of the environment, allowing to predict state transitions and rewards, e.g., giving the current algorithm, the agent could plan future actions and their rewards.

Our solution considers an *agent* that has the goal of generating correct and efficient RBcast algorithms, i.e., algorithms that satisfy the RBcast properties (cf. Section 2.2) and minimizes the efficiency metrics (cf. Section 2.3.4). The solution is composed of a main agent, the *RB-Learner*, that collaborates with an auxiliary agent, the *RB-Oracle*, as represented in Figure 1. The entire execution of the solution for one algorithm is designated a *simulation*. The process starts with the definition of the *inputs* of the learning process. Table 1 summarizes these inputs. The *generation process inputs* include the number of simulations and episodes to run and the rewards and heuristics to be used, the last two containing domain knowledge about the problem (see Sections 4.3 and 4.5). The *validation process inputs* include the specifications to validate the algorithm, such as the failure modes and their tolerance ratios, the number of nodes to model, the properties and the event handlers to be validated. Then, the execution starts with the RB-Learner *generation process*. After that, the RB-Learner gives the generated algorithm to the RB-Oracle, executing the *validation process*, i.e., to assess whether the generated algorithm actually solves the problem. Then, the RB-Oracle returns the validation result to the RB-Learner, which it uses as part of the learning process. These iterations RB-Learner $\rightarrow$ RB-Oracle

Table 2: Actions available to the RB-Learner.

Logic	Message type	Condition threshold	Message type
SEND to all(<t,m>)	*type0*, *type1*	$0, 1, F+1, (N+F)/2, N-F$	*type0*, *type1*
SEND to neighbours(<t,m>)	*type0*, *type1*	$0, 1, F+1, (N+F)/2, N-F$	*type0*, *type1*
SEND to myself(<t,m>)	*type0*, *type1*	$0, 1, F+1, (N+F)/2, N-F$	*type0*, *type1*
DELIVER(m)	-	$0, 1, F+1, (N+F)/2, N-F$	*type0*, *type1*
STOP	-	0	*type0*

→ RB-Learner → ... are repeated many times. Each iteration is designated an *episode*. In the end, the RB-Learner outputs the results of the execution: the most efficient algorithm generated.

With this technique, the idea is that the agent will be able to learn/generate a correct and efficient algorithm by generating multiple algorithms – either correct or incorrect – without the need of prior knowledge of the state of the art in RBcast.

4   RB-Learner

The RB-Learner agent uses Reinforcement Learning to learn not only an algorithm that solves the problem, but also an algorithm that is efficient. Next, we explain the elements behind the learning process of the RB-Learner.

4.1   Actions

An algorithm is composed of a set of event handling routines – two in the case of RBcast, each of which contains a sequence of actions. The RB-Learner selects an *action* from the set of possible actions to add to one of the routines. Each action has two components: the *logic* – the part of the action that is executed – and the *condition* – a statement that must be true in order for the logic to be executed. For example, in action

<p style="text-align:center">SEND to all(<type1,m>) if received (<type0,m>) from 1 distinct process</p>

the logic component is the left-hand part (in red) and the condition is the right-hand part (in blue, starting with the word "if"). Next, we present the logic and condition components that we assume in the paper.

4.1.1   Logic

We selected the following logic components taken from previous works that solve the RBcast problem [9, 26, 29]:

- SEND to all(<t,m>): sends the message *<t,m>* to all processes of the system (including itself);
- SEND to neighbors(<t,m>): sends the message *<t,m>* to all the processes of the system (excluding itself);
- SEND to myself(<t,m>): sends the message *<t,m>* only to itself;
- DELIVER(m): delivers the message *m*;
- STOP: end the execution of the event handler.

In addition to these, we could think of a generic component to send a message to $N - X$ processes, for any $X > 0$ and $X \in \mathbb{N}$. However, this generality is found only in probabilistic algorithms [12, 22], which are outside the scope of this work.

4.1.2   Conditions

In this work, each action is associated with a specific *condition* that defines when the action is executed, and each condition is associated with a specific *threshold*. From a range of works analyzed [9, 26, 29, 52], we selected five thresholds: waiting for *0* (always true, a tautology), 1, $F + 1$, $(N + F)/2$, and $N - F$ messages from different processes. We assume that two conditions are equal if they wait for the same message type and have the same threshold.

4.1.3   Message type

As explained above, messages contain a type: *type0*, *type1*, and so forth. One parameter of the problem of generating an algorithm is how many types of messages it uses. Clearly, the minimum number of types found in algorithms in

the literature corresponds to the maximum number of types needed to solve the problem. For RBcast we found that the number is of two types [29]. Both SEND actions and conditions are influenced by a certain type (except conditions with $threshold = 0$ that wait for no messages).

Table 2 presents all actions available to our agent, resulting in a total of $T = 64$ possible actions – the total number of possible combinations using the components of the table. All actions are associated with all possible conditions except STOP that does not depend on messages being received (or, equivalently, it depends only on the tautology).

We also assume that each action contains an implicit condition that forbids the action to execute more than once for the same message. This means that if a process delivers a message with content $m = 1$, it will never deliver the same message again. The same applies to SEND actions. This inner condition is introduced by previous articles that present RBcast algorithms: for example, in [29] we have the conditions *not yet broadcast* or *not yet RB_delivered* and in [26] we have the condition *if p has not previously executed deliver(R,m)*.

4.2   States

A typical Reinforcement Learning agent interacts with an external *environment*. In our case, the environment is not external, but internal memory. This memory stores the actions already selected by the agent to form the algorithm. Specifically, a state is the sequence of actions selected by the agent up to that moment. By following this representation, the agent will be able to learn to select the best actions based on the actions that the algorithm already contains. Each state follows the algorithm structure defined in Section 2.3.1, being composed by two event handlers and expressed as *State([])*. We assume that *State* A and *State* B are equal if *both contain the same actions, in the same number, in the same event handlers*. The order of the messages inside an event handler is not relevant.

4.3   Rewards

Rewards are used by the agent to learn which actions are suitable or not in each state – a technique called *reward shaping*. In this work, the rewards are related with the *efficiency* (cf. Section 2.3.4) and *correctness* (cf. Section 2.2) of the algorithm: the most efficient correct algorithm will generate the best reward. Rewards are defined as part of the input, but it is important to mention that the agent generates the most efficient algorithm not because of the absolute reward values that we have defined, but because of the relative values between them (the absolute rewards are the result of testing multiple possibilities).

The RB-Agent receives a reward in two moments: (1) every time the agent selects an action – *runtime reward* – and (2) when the agent receives the validation result from the RB-Oracle – *bonus reward*.

4.3.1   Runtime rewards

Runtime rewards are related with the *efficiency* of the produced algorithm, i.e., more efficient algorithms will generate the best rewards. Table 3 summarizes the values we empirically established for calculating these rewards.

The SEND actions and the DELIVER action have a negative reward, as processes need to spend time and resources to execute these actions, so there is a cost involved. For the SEND actions, the SEND to myself action has a better reward than the SEND to neighbours and the SEND to all actions. This happens due to the number of sent messages metric: SEND to myself only sends 1 message, whereas the others involve sending $N - 1$ and N messages, respectively. The threshold of the conditions also influences the reward. Table 3 shows the rewards associated with each selected threshold (c.f Section 4.1.2), following the idea that $0 < 1 \leq F + 1 \leq (N + F)/2 \leq N - F$, where N is the total number of processes and F the maximum number of faulty processes.

Beyond what is presented in Table 3, the addition of a new message type to the algorithm also involves a (negative) reward. Specifically, the reward is added when a SEND action introduces a new type. Each new type has an increasing negative reward: *type0* is associated with the reward 0, *type1* is associated with reward -1, etc. The objective is for the agent to add the minimum number of new types to the algorithm, as each new type involves more communication.

The last aspect that influences the reward obtained by the agent is the event handler where the action is selected. We defined that each action selected for the *RB-Broadcast* event handler has an additional reward of 0, while the actions selected for the *receive* event handler have an additional reward of -1. These rewards favor the addition of actions to the *RB-Broadcast* event handler instead of to the *receive* event handler; this bias is needed because *RB-Broadcast* is executed only once per execution of the algorithm, whereas *receive* is executed N times (one per process), meaning that an action on the *receive* event handler will have a greater impact on the efficiency of the algorithm when compared to an action on the *RB-Broadcast* event handler, e.g, a SEND to myself action will have a cost of 1 message if executed on the *RB-Broadcast* event handler, but a cost of $N * 1$ if executed on the *receive* event handler.

Table 3: Rewards given to each action and each condition.

Logic	Reward	Threshold	Reward
SEND to myself(<t,m>)	−1	0	0
SEND to neighbours(<t,m>)	−2	1	−1
SEND to all(<t,m>)	−3	$F+1$	−2
DELIVER(<m>)	−1	$(N+F)/2$	−3
STOP	0	$N-F$	−4

To summarize, consider the example where the agent chooses the action:

$$\text{SEND to all}(<type0,m>) \text{ if received } (<type0,m>)$$
$$\text{from } N-F \text{ distinct parties}$$

The reward for this action will be -3 (the SEND to all logic) $+\ 0$ (*type0* sent) -4 (the $N-F$ threshold) $=-7$. Then, if the action is selected for the *RB-Broadcast* event handler, it will receive an additional reward of 0 (still a total of -7), while if selected for the *receive* event handler, it will receive an additional reward of -1 (total of -8).

4.3.2   Bonus rewards

After the algorithm is validated by the RB-Oracle, the RB-Learner receives the validation result from the RB-Oracle. The RB-Learner will use that result – correct or incorrect – to get a bonus reward or not. In case the algorithm is correct, there is a bonus of 100, from where we discount the *runtime rewards* accumulated during the generation. For example, if the agent generates a correct algorithm with a runtime reward accumulated of -14 during the state transitions of the generation process, the bonus will be $100+(-14)=86$. This allows the agent to receive a better bonus for the most efficient algorithms. In the case of an incorrect algorithm, the reward received by the agent will be -1: the number of incorrect algorithms will tend to be greater than the correct ones, so we do not want the agent to be severely penalized by finding an incorrect algorithm, since some actions of an incorrect algorithm can still lead to a correct algorithm.

4.4   Generation Process

This section explains one episode of the generation process that generates one algorithm.

The generation process is composed of two development phases: the phase of the *RB-Broadcast* event handler and the phase of the *receive* event handler. Both phases are based on *Q-Learning* [55], one of the most adopted Reinforcement Learning algorithms. This algorithm uses a table designated *QTable* to map the values of each action to each state. Furthermore, each development phase is divided in four steps: (1) heuristic analysis; (2) action selection; (3) reward feedback; and (4) learning update. Next, we explain how the development phases generate the entire algorithm during the generation process.

The generation process begins with the development phase of the RB-broadcast event handler. The agent starts with the internal state empty: *State([])*. Then, comes a loop-based on the current internal state, where the agent applies a set of heuristics (a topic we defer for Section 4.5) to discard the actions not suitable for the current state, from the set of all possible actions (step 1). Then, the agent selects one of the suitable actions based on a policy (step 2). In this work, the agent follows the *Upper Confidence Bound* (UCB) policy [54], a policy based on the idea of *optimistic under uncertainty*. We selected this policy because it allows the agent to find a balance between actions less frequently chosen and actions with higher value, solving the exploration/exploitation problem [16]. Then, based on the action selected, the agent receives a runtime reward (step 3) that it uses to update its learning base, the Qtable, by associating the reward received with the action select on the current state. Moreover, the selected action is added to the current internal state and to the current event handler – the *RB-Broadcast*, originating a new state. For example, if the agent is in the state *State([])* and chooses action A, the action is added to the algorithm and the *RB-Broadcast* event handler, originating the new state *State([action A])*. Finally, the agent defines its current state as the new state and re-executes the development phase, returning to step 1.

The agent continues to re-execute the development phase of the *RB-Broadcast* event handler until the moment when it chooses the STOP action in step 2. By choosing this action, the development phase of the *RB-Broadcast* event handler is completed and the development phase of the *receive* event handler is started. The agent executes this second development phase, but now adds the actions to the *receive* event handler, until the moment when it chooses the STOP action again. This marks the end of the development phase of the *receive* event handler and the completion of the algorithm. The process used in the second development phase (receive) is the same as in the first (broadcast).

Table 4: Heuristics used on the generation process.

Heuristic	Definition
GH1	Do not allow repeated actions on the algorithm
GH2	Allow to define the actions available in each event handler
GH3	Allow to define the conditions available in each event handler
GH4	We can only use a SEND action for each type of message sent and condition
GH5	Messages sent on the *RB-Broadcast* event handler must be of type *type0*
GH6	Allow to define the minimum and maximum number of actions that each event handler can have
GH7	Only select an action that waits for a message type already sent on the algorithm
GH8	The algorithms generated must contain, at least one DELIVER action
GH9	The incorrect states are blocked, in order to never explore them again
GH10	Allows to define the maximum number of types that the algorithm can contain

With the algorithm generated, the generation process ends, and the RB-Learner gives the algorithm to the RB-Oracle that, in turn, will validate it. After validating the algorithm, the RB-Oracle returns the validation result to the RB-Learner that receives a bonus, based on the algorithm being correct or not. This bonus reward will also be used to update the QTable of the agent.

In the first episodes, the generation process will produce random algorithms, led by the policy that allows to *explore* new actions and states. As the simulation progresses, based on the values of the QTable, the policy used by the agent will lead it to *exploit* the actions that have the best value on each state, allowing it to converge to the most efficient algorithms. We use the terms explore and exploit with the precise meanings they have in Reinforcement Learning: explore is related with the search for new and unfamiliar states, whereas exploit refers to the examination of familiar states [16].

4.5 Heuristics

The number of possible algorithms grows exponentially with the base given by the number of actions T (a constant), i.e., has complexity $O(T^i)$, where i is the number of actions. Although the agent has T actions to choose from, there are clearly some bad choices in some cases. For example, it is a bad option to choose STOP as the first action since that would generate an empty algorithm.

To reduce the explosion of possibilities and guide the agent during the generation process, we define a set of *heuristics* [41, 47] for the agent to avoid bad choices. Notice that the heuristics do not help obtaining correct and efficient algorithms; they only reduce the number of possibilities to explore by discarding invalid actions in specific states and, consequently, reducing the time required to find solutions.

Table 4 presents the generation heuristics (GH) that we use to guide the agent in the case of RBcast. These heuristics were defined on the basis of a logic of discarding undesirable actions. Every time the agent is in a state, the agent uses the heuristics to know which actions are available in this specific state, thus to reduce the options from T to $T_h < T$.

GH1 says that the entire algorithm cannot contain duplicate actions (except the STOP action). GH2 allows to define the actions available in each event handler. For RBcast, we define that the agent can select all actions in both event handlers, except the DELIVER action on the *RB-Broadcast* event handler; as the *RB-Broadcast* event handler is only executed by one process, the DELIVER action must exist on the *receive* event handler, so that all processes can deliver the message, thus the possible existence of the DELIVER action on the *RB-Broadcast* event handler is redundant. GH3 allows to define the conditions available in each event handler. Based on this heuristic, we define that on the *receive* event handler, all considered conditions are allowed (see Section 4.1.2). In the *RB-Broadcast*, we only allow conditions based on condition 0, since in that event handler the processes do not receive any message. GH4 allows to define that, for each condition and message type sent, the agent must choose between sending to all, to the neighbours or only to itself. GH5 allows to define a message type for the first communication step of the algorithm (*RB-Broadcast* event handler). GH6 allows to restrict the size of the algorithm generated in terms of the number of actions for each event handler. As previously explained, we took inspiration in one of the most efficient RBcast algorithms [29], so we defined a minimum number of 2 and a maximum number of 4 actions in each event handler. GH7 forces to select actions based on conditions that can be validated, e.g., we forbid the agent to select actions that wait for message types not yet contained in the algorithm. GH8 forces the generation of algorithms with at least one DELIVER action, as that action clearly must exist in the algorithm. GH9 allows to decrease the convergence time by discarding incorrect algorithms that are not related to the solution, which is equivalent to give an infinite negative reward. GH10 allows to define the maximum number of types that the algorithm can contain – in this work, we defined only two possible types (cf. Section 4.1.3).

Table 5: Number of lines of code and programming language used.

Module	Lines of Code	Programming Language
RB-Learner	4216	Python3
Utils	1005	Python3
RB-Oracle	573	Python3 and PROMELA
config	150	Python3 and JSON

5   RB-Oracle

The RB-Oracle is the agent responsible for validating the algorithms generated by the RB-Learner agent, i.e., to implement the *validation process*. This section explains the validation process executed in the context of one episode.

The validation process is responsible for assessing the correctness of the algorithms generated, i.e., for assessing if each algorithm satisfies the RBcast properties (Section 2.2) within one of the variants of the system model (Section 2.1). Every episode, the RB-Learner generates an algorithm and the RB-Oracle validates it.

Automatic validation of a fault-tolerant distributed algorithm can be achieved using different techniques such as *model checking* [23, 31] or *theorem proving* [51]. In this paper, we use a model checking tool called *Spin* [27], a widely used framework on the validation of fault-tolerant algorithms [20, 23, 31, 44] that allows to build models and validate them.

Spin supports a few modes. We use Spin in *simulation* mode, i.e., we use it to simulate the execution of the generated algorithm in a specific system model, doing an exhaustive exploration of the state space. In essence, during the state space exploration, Spin verifies if none of the three RBcast properties (RB-Agreement, RB-Validity, and RB-Integrity) is violated. The three properties, the protocol, the values of N and F, the system architecture, the behavior of each failure mode (No/Crash/Byzantine-Failure mode), the process that initiates the verification (randomly selected), and the faulty processes (also randomly selected, for $F > 0$) are all specified in PROMELA (Process or Protocol Meta Language).

The RB-Oracle validates algorithms considering three failure modes: No-Failure, Crash-Failure, and Byzantine-Failure.

For the *No-Failure mode*, the RB-Oracle verifies the algorithm considering that all processes are correct, i.e., by following the actions of the algorithm without deviations. In this mode, in the experiments, we assume a system with $N = 3$ processes and $F = 0$. Moreover, we model only one possible verification of the system: since the algorithm that runs in each process is the same, more than one model would be redundant.

In the *Crash-Failure mode*, we simulate the crash of the process assuming the worst case possible: crash failures happen between the sending of messages, since the impact of a crash failure is the highest when it leads a message to be delivered only to a fraction of the processes. In this mode, in the experiments, we assume a system with $F = 1$ faulty processes and $N = 3$ processes, the minimum necessary to have $F = 1$ failures (see Section 2.1). Moreover, for this mode we build two models: when the process that initiates the algorithm is correct and when it is faulty. This allows verifying cases when either of the event handlers fail.

For the *Byzantine-Failure mode*, the RB-Oracle models a range of attacks where all faulty processes send the same malicious message to a predefined group of correct processes – from a group with 0 processes, and consequently, not sending to anyone, to sending to $N - F$ processes, and consequently, sending to all correct processes. In this mode, in the experiments, we assume a system with $F = 1$ faulty processes and $N = 4$ processes, which is the minimum number of processes necessary to have $F = 1$ faulty processes (see Section 2.1). Moreover, similar to the process on the Crash-Failure mode, in this mode we also build two models: one to model a failure on each event handler of the algorithm.

6   Implementation

As explained in earlier sections, our approach is based on two agents. Table 5 shows the number of lines of code and the programming language used to implement the entire solution. The Utils module corresponds to the functions used by both the RB-Learner and RB-Oracle agents and the config module to the configuration files used as inputs for our system.

The RB-Learner uses the Q-Learning algorithm and the UCB policy to implement the generation and learning of the algorithms. The RB-Oracle models a distributed system, with N processes and F faulty processes, simulating the execution of the algorithm generated using the Spin framework. After receiving the algorithm from the RB-Learner,

the RB-Oracle creates a validation model and stores it in a PROMELA language file (*.pml* extension). Then, based on the *.pml* file, the RB-Oracle generates a verification file (*pan.c*) – a *C* program that performs a verification of the correctness requirements for the system – and compiles it using *gcc*, generating an executable file. Lastly, the agent uses Spin to run the executable file and, with that, check the correctness of the algorithm.

The normal functioning of our solution is, as introduced in Section 3, organized in *episodes* and *simulations*: one episode is composed by one generation process and one validation process, while one simulation corresponds to the entire execution with usually many episodes.

7 Experimental Evaluation

Our experimental evaluation aims to assess the effectiveness and correctness of our solution. It will answer the following questions:

1. How many states does the agent explore until finding the first correct algorithm and converging to the most efficient algorithm?
2. How many algorithms are generated in total, in each experiment?
3. How many algorithms are generated until finding the first correct algorithm?
4. What is the proportion of correct and incorrect algorithms, from the total number of generated algorithms?
5. How does each proposed heuristic influence the learning process?

Table 6: Experimental evaluation inputs.

Generation process inputs	Value
Number of simulations	5
Number of episodes	12.000
Rewards	Defined at Table 3
Heuristics to be applied and configurations associated	Configured as presented on Section 4.5
Validation process inputs	**Value**
Failure modes to test	No-Failure, Crash-Failure and Byzantine-Failure modes
Number N of nodes to model	4
Fault tolerance ratios	$F = 0$ for the No-Failure, $F = \lfloor (N-1)/2 \rfloor$ for the Crash-Failure and $F = \lfloor (N-1)/3 \rfloor$ for the Byzantine-Failure
Properties to be validated	*RB-Agreement*, *RB-Validity* and *RB-Integrity*
Events where faults can occur	Both *RB-Broadcast* and *receive* events

We ran our experiments on a single machine with 32 vCPUs, 64 GB of memory and Debian 10. We made experiments for three cases: (1) No-Failure mode only; (2) No-Failure and Crash-Failure modes; and (3) No-Failure, Crash-Failure and Byzantine-Failure modes. Each experiment is referred to by the last failure mode of the list (the most generic): No-Failure, Crash-Failure and Byzantine-Failure experiments. All results shown in the next sections are, except when noticed, averages of 5 simulations runs, each with $12,000$ episodes – the minimum number of episodes that we have found to be possible for the agent to converge to the most efficient algorithms in all experiments. The $12,000$ episodes took ± 9 hours to run on the No-Failure experiment, ± 3 days to run on the Crash-Failure experiment and ± 7 days to run on the Byzantine-Failure experiment. This increase in time is due to the time needed for Spin to verify the models. Table 6 summarize the inputs considered for the experimental evaluation.

7.1 States Explored

For each algorithm generated, the agent explores multiple states when selecting the actions. This first set of experiments assesses the number of states explored in each experiment. Figure 2 shows the total number of states explored by the agent for each episode, on the entire experiment. Table 7 shows the number of states explored until the agent generated the first correct algorithm.

As expected, we can see in both figures that the agent needs to explore more states when the complexity of the problem to solve increases, i.e., the agent needs to explore more states when trying to find a Byzantine-tolerant algorithm – almost $20,000$ states – when compared to a Crash-tolerant or a non-fault-tolerant algorithm – around $12,000$ and
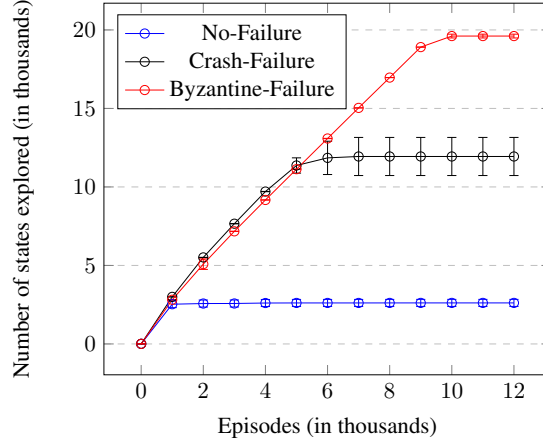
Figure 2: Number of states explored during each experiment.

Table 7: Number of states generated until the first correct algorithm is generated.

Experiment	States
No-Failure	15.4 ± 0.0
Crash-Failure	33.4 ± 16.9
Byzantine-Failure	13462.8 ± 6909.4

$3,000$ states, respectively. Additionally, the agent also takes more time to converge when trying to find a Byzantine-tolerant algorithm – between $8,000$ and $10,000$ episodes – when compared to the other cases – between $1,000$ and $2,000$ episodes for the No-Failure algorithm and between $4,000$ and $6,000$ episodes for the Crash-Failure algorithm.

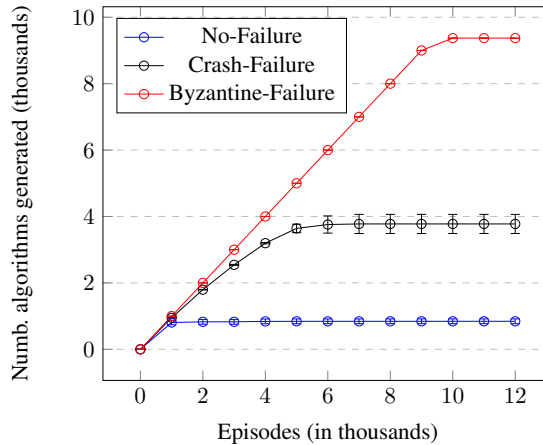## 7.2 Algorithms Generated

Figure 3: Number of algorithms generated during each experiment.

The agent generates multiple algorithms with the objective to learn from them and, therefore, we decided to assess how many algorithms the agent generates in each experiment. Figure 3 shows the number of algorithms generated by the agent per episode. Table 8 shows the number of correct and incorrect algorithms generated, as also the number of algorithms generated until the first correct algorithm. As expected, and similarly to what happens with the number of states, the agent needs to generate more algorithms, as also takes more time to converge, with the increase of the complexity of the problem to solve. Moreover, another interesting aspect is the percentage of incorrect algorithms generated by the agent in each test: $\pm 60\%$ on the No-Failure test, $\pm 89\%$ on the Crash-Failure test and 99.9% on the

Table 8: Number of correct and incorrect algorithms generated in each experiment and number of algorithms generated until the first correct algorithm.

Experiment	Correct Algo.	Incorrect Algo.	Number Algo.
No-Fail.	341.8 ± 23.2	502.0 ± 53.3	2.6 ± 1.9
Crash-Fail.	420.2 ± 32.1	3355.6 ± 256.6	6.8 ± 4.1
Byzantine-Fail.	2.0 ± 0.0	9369.6 ± 6.8	6260.0 ± 3467.9

Byzantine-Failure test, which means that, even with all the Heuristics defined, the agent still has a difficult task to generate a correct algorithm.

The final algorithms generated by each experiment are presented in Algorithms 1, 2 and 3, one for each failure mode.

Algorithm 1 Most efficient RBcast algorithm for a No-Failure experiment generated by the RB-Learner.

1: *when RB-Broadcast(m) do*:
2:    SEND to all($<type0$,m$>$) if received ($<type0$,m$>$) from 0 distinct parties and not already sent;
3:    STOP if received ($<type0$,m$>$) from 0 distinct parties;

4: *when receive(m) do*:
5:    DELIVER($<$m$>$) if received ($<type0$,m$>$) from 0 distinct parties and not already delivered;
6:    STOP if received ($<type0$,m$>$) from 0 distinct parties;

In the No-Failure mode, the agent converged to Algorithm 1 in 4 of the simulations executed. This algorithm is equivalent to one presented in [12]: both exchange, at most, N messages, require 1 communication step and 1 message type (or none) and need to receive 1 message.

On the Crash-Failure mode, the agent also converged to Algorithm 2 in 4 of the simulations executed. Note that the algorithm sends a new message type on the *receive* event handler (*type1*), when it could send the *type0*. This happens because of the heuristic GH5, that only allows to send messages of *type0* on the *RB-Broadcast* event handler. This algorithm is similar to the one presented in [26, 52]: both exchange, at most, $N^2 - N + 1$ messages, require 1 communication step and 1 message type (or none) and need to receive 1 message.

On the Byzantine-Failure mode, the agent converged to Algorithm 3 in all the simulations executed. This algorithm is one of the most efficient algorithms developed and it is similar to the one presented in [29]: both exchange, at most, $N^2 + N$ messages, require 2 communication steps and 2 message types and need to receive $(N + F)/2$ messages.

7.3   Impact of the Heuristics

The heuristics we defined (see Section 4.5) guide the agent by helping it to avoid incorrect algorithms and reduce the number of states to explore. In this evaluation, we analyzed the importance of each heuristic with the Crash-Failure experiment.

To achieve this, we ran one experiment with each GH turned off and all others turned on. There were two exceptions. In GH6, we increased the maximum number of actions in each event from 4 to 5, but did not turn this heuristic off, to avoid that the agent would generate algorithms with too many actions. For GH10, we increased the maximum number of types from 2 to 3 but did not turn it off, as the agent could explore too many types. We executed one simulation with 10, 000 episodes for each experiment.

Figure 4 shows the evolution of the number of algorithms generated with each GH turned off, from where we conclude that all heuristics are important to reduce the number of states explored until a correct and efficient algorithm is obtained.

8   Related Work

Fault-tolerant algorithms have been widely studied over the years [2, 7–9, 14, 15, 18, 19, 29, 30, 52]. These algorithms: solve different problems, such as *Reliable Broadcast* [29] and *Consensus* [14]; tolerate different failure modes, like *Crash* [26] and *Byzantine* [19]; use different communication models, namely *fully-connected* [18] and *partially-connected* [8]; and tolerate different fault ratios, such as $\lfloor (N-1)/3 \rfloor$ [9] and $\lfloor (N-1)/2 \rfloor$ [19], where N is the number of components in the system. However, as far as we know, all works are based on manual or brute-force processes, without any kind of artificial intelligence helping with the process of generating the algorithms.

Algorithm 2 Most efficient RBcast algorithm for a Crash-Failure experiment generated by the RB-Learner.

1: *when RB-Broadcast(m) do*:
2:   SEND to myself(*<type0*,m>) if received (*<type0*,m>) from 0 distinct parties and not already sent;
3:   STOP if received (*<type0*,m>) from 0 distinct parties;

4: *when receive(m) do*:
5:   SEND to neighbours(*<type1*,m>) if received (*<type0*,m>) from 0 distinct parties and not already sent;
6:   DELIVER(<m>) if received (*<type0*,m>) from 0 distinct parties and not already delivered;
7:   STOP if received (*<type0*,m>) from 0 distinct parties;

Algorithm 3 Most efficient RBcast algorithm for a Byzantine-Failure experiment generated by the RB-Learner.

1: *when RB-Broadcast(m) do*:
2:   SEND to all(*<type0*,m>) if received (*<type0*,m>) from 0 distinct parties and not already sent;
3:   STOP if received (*<type0*,m>) from 0 distinct parties;

4: *when receive(m) do*:
5:   SEND to all(*<type1*,m>) if received (*<type0*,m>) from 1 distinct party and not already sent;
6:   DELIVER(<m>) if received (*<type1*,m>) from $(N + F)/2$ distinct parties and not already delivered;
7:   SEND to all(*<type1*,m>) if received (*<type1*,m>) from $F + 1$ distinct parties and not already sent;
8:   STOP if received (*<type0*,m>) from 0 distinct parties;

Automatic code generation works focused on local, non-distributed, single-threaded code, started by using static techniques such as design-based approaches [11], UML [46] or reverse engineering [49], but lately has been focused on using machine learning techniques [1,4,57] mostly supervised learning techniques, such as Deep Learning [40,58]. For distributed code, we identified two works: one that automatically finds mutual exclusion algorithms [5] and another that automatically investigates and validates Consensus algorithms [59], both using brute-force approaches. However, in our case, we use an unsupervised machine learning technique – Reinforcement Learning [32,54,56] – that allows the agent to learn without the need of prior knowledge about solutions to the problem to be solved. Reinforcement Learning has been explored mainly in games [37,45,53] and robotics [28,34,43], so our work applies it to an entirely different problem. When compared to our approach, the most interesting works that use this technique are those that use an agent to generate GPU compiler heuristics [17], an agent that chooses the most suitable algorithm [36], or an agent that is capable of generating experiment input data [33]. Nevertheless, these problems are very different from ours.

For the validation of fault-tolerant algorithms [23,25,38,51], we have identified and used the Spin/PROMELA [27] model checker [23,31] framework, since it allows us to model and validate the generated algorithms, is used by a significant number of related works [20,23,31,44] and also has good community support resources. Besides
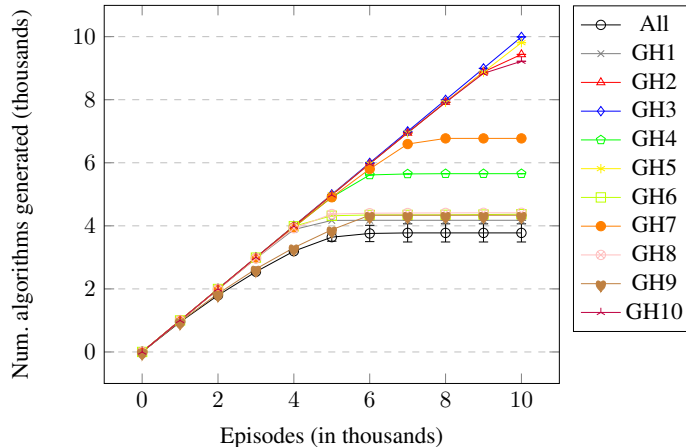

Figure 4: Number of algorithms generated during each experiment without the identified GH. The All line represents a experiment with all GH turned on.

13

Spin/PROMELA, other possible languages and frameworks could have been used: TLA+[2] [38], the ByMC[3] framework [35] and the IC3PO[4] [24].

9 Conclusion

Fault-tolerant algorithms have been studied over the years, discussing different problems and variants. However, this study is complex and was always based on a human-oriented process.

To automate this process, we propose a solution based on two agents, RB-Learner and RB-Oracle, capable of learning to generate a distributed algorithm, more precisely the RBcast algorithm. As we have presented during the experimental evaluation, our solution is capable of generating correct and efficient algorithms, depending on the properties of the problem, which proves that our approach can be used to generate distributed algorithms.

To the best of our knowledge, this work is the first that merges both areas of generation and validation into an automatic process capable of generating correct and efficient RBcast algorithms. Additionally, this is the first work that uses a machine learning approach to generate correct and efficient algorithms to solve a specific distributed problem.

For further research, we aim to apply our approach to different distributed problems, e.g. Consensus, and try to decrease the number of inputs needed, to further decouple our agent from knowledge based on previous works, e.g. the threshold expressions.

References

[1] Miltiadis Allamanis, Earl T Barr, Premkumar Devanbu, and Charles Sutton. A survey of machine learning for big code and naturalness. *ACM Computing Surveys*, 51(4):1–37, 2018.

[2] Tim Anema. Message efficient Byzantine reliable broadcast protocols on known topologies. Bachelor's Thesis, TU Delft, 2021.

[3] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, et al. A view of cloud computing. *Communications of the ACM*, 53(4):50–58, 2010.

[4] Batuhan Aşıroğlu, Büşta Rümeysa Mete, Eyyüp Yıldız, Yağız Nalçakan, Alper Sezen, Mustafa Dağtekin, and Tolga Ensari. Automatic HTML code generation from mock-up images using machine learning techniques. In *2019 Scientific Meeting on Electrical-Electronics & Biomedical Engineering and Computer Science*, pages 1–4, 2019.

[5] Yoah Bar-David and Gadi Taubenfeld. Automatic discovery of mutual exclusion algorithms. In *International Symposium on Distributed Computing*, pages 136–150. Springer, 2003.

[6] Rida A Bazzi and Gil Neiger. Simplifying fault-tolerance: providing the abstraction of crash failures. *Journal of the ACM*, 48(3):499–554, 2001.

[7] Michael Ben-Or. Another advantage of free choice: Completely asynchronous agreement protocols. In *Proceedings of the 2nd Annual ACM Symposium on Principles of Distributed Computing*, pages 27–30, 1983.

[8] Silvia Bonomi, Jérémie Decouchant, Giovanni Farina, Vincent Rahli, and Sébastien Tixeuil. Practical Byzantine reliable broadcast on partially connected networks. *arXiv preprint arXiv:2104.03673*, 2021.

[9] Gabriel Bracha. An asynchronous [(n-1)/3]-resilient consensus protocol. In *Proceedings of the 3rd Annual ACM Symposium on Principles of Distributed Computing*, pages 154–162, 1984.

[10] Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM*, 32(4):824–840, 1985.

[11] Frank J. Budinsky, Marilyn A. Finnie, John M. Vlissides, and Patsy S. Yu. Automatic code generation from design patterns. *IBM Systems Journal*, 35(2):151–171, 1996.

[12] Christian Cachin, Rachid Guerraoui, and Luís Rodrigues. *Introduction to reliable and secure distributed programming*. Springer Science & Business Media, 2011.

[13] Christian Cachin, Klaus Kursawe, Frank Petzold, and Victor Shoup. Secure and efficient asynchronous broadcast protocols. In *Annual International Cryptology Conference*, pages 524–541. Springer, 2001.

---

[2] https://lamport.azurewebsites.net/tla/tla.html

[3] https://github.com/konnov/bymc

[4] https://github.com/aman-goel/ic3po

[14] Miguel Castro and Barbara Liskov. Practical Byzantine fault tolerance. In *Proceedings of the 3rd USENIX Symposium on Operating Systems Design and Implementation*, pages 173–186, 1999.

[15] Jo-Mei Chang and Nicholas F. Maxemchuk. Reliable broadcast protocols. *ACM Transactions on Computer Systems*, 2(3):251–273, 1984.

[16] Melanie Coggan. Exploration and exploitation in reinforcement learning. *Research supervised by Prof. Doina Precup, CRA-W DMP Project at McGill University*, 2004.

[17] Ian Colbert, Jake Daly, and Norm Rubin. Generating GPU compiler heuristics using reinforcement learning. *arXiv preprint arXiv:2111.12055*, 2021.

[18] Miguel Correia, Nuno Ferreira Neves, and Paulo Veríssimo. From consensus to atomic broadcast: Time-free Byzantine-resistant protocols without signatures. *The Computer Journal*, 49(1):82–96, 2006.

[19] Miguel Correia, Giuliana S Veronese, and Lau Cheuk Lung. Asynchronous Byzantine consensus with 2f+1 processes. In *Proceedings of the 2010 ACM Symposium on Applied Computing*, pages 475–480, 2010.

[20] Giorgio Delzanno, Michele Tatarek, and Riccardo Traverso. Model checking paxos in spin. *Electronic Proceedings in Theoretical Computer Science*, 161:131–146, aug 2014.

[21] Cynthia Dwork and Yoram Moses. Knowledge and common knowledge in a Byzantine environment: crash failures. *Information and Computation*, 88(2):156–186, 1990.

[22] Patrick Th Eugster, Rachid Guerraoui, Sidath B Handurukande, Petr Kouznetsov, and A-M Kermarrec. Lightweight probabilistic broadcast. *ACM Transactions on Computer Systems*, 21(4):341–374, 2003.

[23] Annu Gmeiner, Igor Konnov, Ulrich Schmid, Helmut Veith, and Josef Widder. Tutorial on parameterized model checking of fault-tolerant distributed algorithms. In *International School on Formal Methods for the Design of Computer, Communication and Software Systems*, pages 122–171. Springer, 2014.

[24] Aman Goel and Karem Sakallah. On symmetry and quantification: A new approach to verify distributed protocols. In *NASA Formal Methods Symposium*, pages 131–150. Springer, 2021.

[25] Aman Goel and Karem A. Sakallah. Towards an automatic proof of Lamport's Paxos. In Ruzica Piskac and Michael W Whalen, editors, *Formal Methods in Computer-Aided Design*, pages 112–122, New Haven, Connecticut, October 2021. doi:https://doi.org/10.34727/2021/isbn.978-3-85448-046-4_20.

[26] Vassos Hadzilacos and Sam Toueg. A modular approach to fault-tolerant broadcasts and related problems. Technical Report TR94-1425, Cornell University, Department of Computer Science, May 1994.

[27] Gerard J. Holzmann. The model checker SPIN. *IEEE Transactions on Software Engineering*, 23(5):279–295, 1997.

[28] Julian Ibarz, Jie Tan, Chelsea Finn, Mrinal Kalakrishnan, Peter Pastor, and Sergey Levine. How to train your robot with deep reinforcement learning: lessons we have learned. *The International Journal of Robotics Research*, 40(4-5):698–721, 2021.

[29] Damien Imbs and Michel Raynal. Simple and efficient reliable broadcast in the presence of Byzantine processes. *arXiv preprint arXiv:1510.06882*, 2015.

[30] Damien Imbs and Michel Raynal. Trading off t-resilience for efficiency in asynchronous Byzantine reliable broadcast. *Parallel Processing Letters*, 26(04):1650017, 2016.

[31] Annu John, Igor Konnov, Ulrich Schmid, Helmut Veith, and Josef Widder. Towards modeling and model checking fault-tolerant distributed algorithms. In *International SPIN Workshop on Model Checking of Software*, pages 209–226. Springer, 2013.

[32] Leslie Pack Kaelbling, Michael L Littman, and Andrew W Moore. Reinforcement learning: A survey. *Journal of Artificial Intelligence Research*, 4:237–285, 1996.

[33] Junhwi Kim, Minhyuk Kwon, and Shin Yoo. Generating test input with deep reinforcement learning. In *2018 IEEE/ACM 11th International Workshop on Search-Based Software Testing*, pages 51–58, 2018.

[34] Jens Kober, J Andrew Bagnell, and Jan Peters. Reinforcement learning in robotics: A survey. *The International Journal of Robotics Research*, 32(11):1238–1274, 2013.

[35] Igor Konnov and Josef Widder. ByMC: Byzantine model checker. In *International Symposium on Leveraging Applications of Formal Methods*, pages 327–342. Springer, 2018.

[36] Michail G. Lagoudakis and Michael L. Littman. Algorithm selection using reinforcement learning. In *Proceedings of the 17th International Conference on Machine Learning*, pages 511–518, 2000.

[37] Guillaume Lample and Devendra Singh Chaplot. Playing FPS games with deep reinforcement learning. In *31st AAAI Conference on Artificial Intelligence*, 2017.

[38] Leslie Lamport and Stephan Merz. Specifying and verifying fault-tolerant systems. In *Formal Techniques in real-time and fault-tolerant systems*, pages 41–76. Springer, 1994.

[39] Leslie Lamport, Robert Shostak, and Marshall Pease. The Byzantine generals problem. *ACM Transactions on Programming Languages and Systems*, 4(3):382–401, 1982.

[40] Triet HM Le, Hao Chen, and Muhammad Ali Babar. Deep learning for source code modeling and generation: Models, applications, and challenges. *ACM Computing Surveys*, 53(3):1–38, 2020.

[41] Douglas B Lenat. The nature of heuristics. *Artificial Intelligence*, 19(2):189–249, 1982.

[42] Nancy A Lynch. *Distributed Algorithms*. Morgan Kaufmann, 1996.

[43] Maja J Matarić. Reinforcement learning in the multi-robot domain. In *Robot Colonies*, pages 73–83. Springer, 1997.

[44] Takahiro Minamikawa, Tatsuhiro Tsuchiya, and Tohru Kikuno. Language and tool support for model checking of fault-tolerant distributed algorithms. In *2008 14th IEEE Pacific Rim International Symposium on Dependable Computing*, pages 40–47, 2008. `doi:10.1109/PRDC.2008.13`.

[45] Volodymyr Mnih, Koray Kavukcuoglu, David Silver, Alex Graves, Ioannis Antonoglou, Daan Wierstra, and Martin Riedmiller. Playing atari with deep reinforcement learning. *arXiv preprint arXiv:1312.5602*, 2013.

[46] Tomas G. Moreira, Marco A. Wehrmeister, Carlos E Pereira, Jean-Francois Petin, and Eric Levrat. Automatic code generation for embedded systems: From UML specifications to VHDL code. In *8th IEEE International Conference on Industrial Informatics*, pages 1085–1090, 2010.

[47] Heiner Müller-Merbach. Heuristics and their design: a survey. *European Journal of Operational Research*, 8(1):1–23, 1981.

[48] P Russel Norvig and S Artificial Intelligence. *A modern approach*. Prentice Hall Upper Saddle River, NJ, USA:, 2002.

[49] Pedro Orvalho, Miguel Terra-Neves, Miguel Ventura, Ruben Martins, and Vasco Manquinho. SQUARES: a SQL synthesizer using query reverse engineering. *Proceedings of the VLDB Endowment*, 13(12):2853–2856, 2020.

[50] Morgen E Peck. Blockchains: How they work and why they'll change the world. *IEEE Spectrum*, 54(10):26–35, 2017.

[51] Vincent Rahli, Ivana Vukotic, Marcus Völp, and Paulo Esteves-Verissimo. Velisarios: Byzantine fault-tolerant protocols powered by coq. In *European Symposium on Programming*, pages 619–650. Springer, 2018.

[52] Michel Raynal. *Fault-Tolerant Message-Passing Distributed Systems*. Springer, 2018.

[53] David Silver, Thomas Hubert, Julian Schrittwieser, Ioannis Antonoglou, Matthew Lai, Arthur Guez, Marc Lanctot, Laurent Sifre, Dharshan Kumaran, Thore Graepel, et al. A general reinforcement learning algorithm that masters chess, shogi, and go through self-play. *Science*, 362(6419):1140–1144, 2018.

[54] Richard S Sutton and Andrew G Barto. *Reinforcement Learning: An Introduction*. MIT Press, 2018.

[55] Christopher JCH Watkins and Peter Dayan. Q-learning. *Machine Learning*, 8(3-4):279–292, 1992.

[56] Marco A Wiering and Martijn Van Otterlo. Reinforcement learning. *Adaptation, Learning, and Optimization*, 12(3), 2012.

[57] Xiaojiang Zhang, Ying Jiang, and Zhijun Wang. Analysis of automatic code generation tools based on machine learning. In *2019 IEEE International Conference on Computer Science and Educational Informatization*, pages 263–270, 2019.

[58] Qihao Zhu and Wenjie Zhang. Code generation based on deep learning: a brief review. *arXiv preprint arXiv:2106.08253*, 2021.

[59] Piotr Zielinski. Automatic verification and discovery of byzantine consensus protocols. In *37th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN'07)*, pages 72–81. IEEE, 2007.