

CGAAL: Distributed On-The-Fly ATL Model Checker with Heuristics*

Falke B. Ø. Carlsen

Department of Computer Science,
Aalborg University, Denmark
falkeboc@cs.aau.dk

Jener Rasmussen

jener@jener.dk

Lars Bo P. Frydenskov

Department of Computer Science,
Aalborg University, Denmark
larsbopark@gmail.com

Mathias M. Sørensen

mathiasmehlsoerenen@gmail.com

Nicolaj Ø. Jensen

Department of Computer Science,
Aalborg University, Denmark
noje@cs.aau.dk

Asger G. Weirsøe

asger@weircon.dk

Mathias C. Jensen

Department of Computer Science,
Aalborg University, Denmark
mcje@cs.aau.dk

Kim G. Larsen

Department of Computer Science,
Aalborg University, Denmark
kgl@cs.aau.dk

We present CGAAL, our efficient on-the-fly model checker for alternating-time temporal logic (ATL) on concurrent game structures (CGS). We present how our tool encodes ATL as extended dependency graphs with negation edges and employs the distributed on-the-fly algorithm by Dalsgaard et al. Our tool offers multiple novel search strategies for the algorithm, including DHS which is inspired by PageRank and uses the in-degree of configurations as a heuristic, IHS which estimates instability of assignment values, and LPS which estimates the distance to a state satisfying the constituent property using linear programming. CGS are input using our modelling language LCGS, where composition and synchronisation are easily described. We prove the correctness of our encoding, and our experiments show that our tool CGAAL is often one to three orders of magnitude faster than the popular tool PRISM-games on case studies from PRISM’s documentation and among case studies we have developed. In our evaluation, we also compare and evaluate our search strategies, and find that our custom search strategies are often significantly faster than the usual breadth-first and depth-first search strategies.

1 Introduction

Software plays a large role in our everyday lives, making decisions, enabling efficient communication, ensuring safety, and many more critical tasks. Furthermore, the complexity of the software and the decisions they have to make are ever-increasing due to the interconnectivity and reactive nature of modern systems. These software systems mutually depend on, communicate with, and guide each other based on their collective and internal states. Even a few interconnected systems that are not necessarily themselves too complex can give rise to an extremely complex system as a whole. Safety-critical software that operates in contexts where errors could lead to human casualties or significant capital losses requires methods to verify their correctness. Unfortunately, such methods are difficult to implement due to the sheer state-space explosion that arises from the inherent parallel composition of systems.

In this paper, we consider systems that can be expressed as discrete multiplayer games in which the actors can perform actions concurrently. That is, in each configuration of the game, each player simultaneously chooses an action they wish to perform and the resulting decision vector then deterministically

*Funded by the VILLUM INVESTIGATOR project S4OS.

results in the next configuration of the game. Such concurrent games are extremely expressive and can be used to model a variety of complex systems and lends themselves inherently, by their concurrent nature, to describing multi-actor systems. The properties that we wish to verify on these systems are those that can be expressed by Alternating-time Temporal Logic (ATL) [1], a game theoretic extension of Computation Tree Logic (CTL) [2, 3], that like CTL can be used to describe safety and liveness properties. ATL extends CTL by replacing the existential and universal path quantifiers with so-called coalition quantifiers. These coalition quantifiers describe properties of the game in which a coalition of players can either enforce some desired outcome or cannot avoid an outcome. Of course, proof of satisfaction of such property is done by finding a strategy showing that the players can indeed enforce said property or by showing that no matter what strategy the players follow they cannot avoid said property.

One existing tool that can model check concurrent (stochastic) games is PRISM-games [12], an extension of the probabilistic model checker PRISM [11]. PRISM excels at stochastic models such as discrete- and continuous-time Markov chains, Markov decision processes, and (priced) probabilistic timed automata, for which it can verify various properties described in LTL, CSL, and probabilistic CTL*. There are multiple verification engines in PRISM, both symbolic and explicit-state. The PRISM-games extension focuses on stochastic games in which game-theoretic approaches are needed. PRISM-games can among other things synthesise strategies and reason about equilibria-based properties, where players may have distinct, but not necessarily conflicting objectives. These properties are typically described with probabilistic ATL with rewards. In version 3.0 [12], PRISM-games added support for concurrent stochastic games.

In [14], Liu and Smolka present a global and a local algorithm to compute fixed-point boolean vertex assignments in dependency graphs with directed hyper-edges representing dependencies between vertex values. The global algorithm has a better worst-case running time of the two, but the on-the-fly local algorithm only explores the graph as needed to compute the fixed-point assignment. Many problems have since been encoded as dependency graphs. Notably, model-checking problems for CTL can be encoded as dependency graphs and the algorithms by Liu and Smolka can be used to compute the satisfaction relation enabling model-checking using dependency graphs. J. F. Jensen et al. [9] show that weighted CTL can be encoded in dependency graphs as well and introduce symbolic edges to handle weights efficiently. In [10], M. C. Jensen et al. introduce symbolic dependency graphs, an encoding of probabilistic weighted CTL, and an implementation of a local and global algorithm. A. Dalsgaard et al. [4] present a distributed extension to the local algorithm to boost the performance and show how to model Petri Nets problems using CTL and dependency graphs. Furthermore, they extend dependency graphs by adding components and negation edges to solve formulae with negations. An abstract dependency graph framework is presented by S. Enevoldsen et al. in [6] where they generalise the various extensions of the dependency graphs with a general algorithm for any Noetherian partial order domain. In the closely related work of [5], S. Enevoldsen et al. uses abstract dependency graphs to model check probabilistic ATL for weighted stochastic games. In multiple works [4, 10] the local algorithm is found to be faster more often than the global algorithm in practice.

In this paper, we introduce the tool CGAAL (Concurrent Games Aalborg), a model checker implemented in Rust that allows for the verification of ATL properties on concurrent games. CGAAL works by encoding the model checking problem as an extended dependency graph and by implementing the on-the-fly local distributed algorithm by Dalsgaard et al. [4]. The concurrent games are described and inputted to the tool using LCGS, our PRISM-language inspired model language, and for specific ATL formulae, we can output a strategy describing how to ensure satisfaction given a positive result. Furthermore, CGAAL implements a variety of strategies for searching the problem state space: a heuristic inspired by PageRank [15] and two heuristic search strategies that exploit the state vector representation

of our model, alongside the usual breadth-first and depth-first approach. We conduct experiments comparing our different search strategies with each other on several case studies and show that having more available compute threads often leads to speed-ups of 1-2 orders of magnitude. We also perform experiments comparing our tool on a selection of case studies to the state-of-the-art PRISM-games. Though while PRISM-games is a tool specialised for probabilistic games we find that our implementation often outperforms PRISM-games and especially whenever we are not required to compute the entire fixed point in which case we are often 1-3 orders of magnitude faster. Again, we emphasise that PRISM-games is meant for probabilistic games and as such this is not a completely fair one-to-one comparison.

Outline. This paper is structured as follows. Section 2 introduces the formal definitions of concurrent games and alternating-time temporal logic. In Section 3 we present how model checking of ATL properties in concurrent games can be encoded in extended dependency graphs. We also give a short explanation of the CERTAINZERO algorithm and our search strategies. In Section 4 we give a short example of how to use CGAAL and the LCGS language. An evaluation of our tool and a comparison with PRISM can be found in Section 5, and we conclude on our findings in Section 6.

2 Definitions

We recall the definitions of concurrent games and alternating-time temporal logic.

Concurrent Games A *concurrent game structure* (CGS) is a tuple $S = \langle k, Q, \Pi, \pi, d, \delta \rangle$ where:

- $k \geq 1$ is a natural number of players. We identify the players with the numbers $1, \dots, k$.
- Q is a finite set of states.
- Π is a finite set of atomic propositions, also called labels.
- A set $\pi(q) \subseteq \Pi$ of propositions are true at $q \in Q$. The function π is called the *labelling function*.
- For each player $a \in \{1, \dots, k\}$ and each state $q \in Q$, a natural number $d_a(q) \geq 1$ of moves are available to player a at state q . The moves of player a at state q are identified with the numbers $1, \dots, d_a(q)$. A move vector at $q \in Q$ is a tuple $v = \langle j_1, \dots, j_k \rangle$ such that $1 \leq j_a \leq d_a(q)$ for each player $a \in \{1, \dots, k\}$. Additionally, given a state $q \in Q$, we write $D(q)$ for the set $\{1, \dots, d_1(q)\} \times \dots \times \{1, \dots, d_k(q)\}$ of all move vectors possible at q . The function D is called the *move function*.
- For each state $q \in Q$ and each move vector $v = \langle j_1, \dots, j_k \rangle \in D(q)$, we have that $\delta(q, v) \in Q$ is the resulting state when each player $a \in \{1, \dots, k\}$ chooses move j_a in the state q . The function δ is called the *transition function*.

A state q' is a *successor* to q if and only if there exists a transition from q to q' , i.e. there exists a move vector $v \in D(q)$ such that $\delta(q, v) = q'$. A *computation* is an infinite sequence $\lambda = q_0, q_1, q_2, \dots$ of states, such that for all $i \geq 0$, the state q_{i+1} is a successor of q_i . Given a computation λ and a position $i \geq 0$, we use the following notations:

| | |
|----------------------|---|
| $\lambda[i]$ | The i th state of the computation of λ |
| $\lambda[0, i]$ | The finite prefix q_0, q_1, \dots, q_i of λ |
| $\lambda[i, \infty]$ | The infinite suffix $q_i, q_{i+1}, q_{i+2}, \dots$ of λ |

A computation starting in the state q is called a q -computation.

Game strategies Consider a concurrent game structure $S = \langle k, Q, \Pi, \pi, d, \delta \rangle$ over the set $\Sigma = \{1, \dots, k\}$ of players. A *strategy* for a player $a \in \Sigma$ is a function $z_a : Q \rightarrow \mathbb{N}$ that maps every state $q \in Q$ to a natural number, such that $z_a(q) \leq d_a(q)$. In other words, the strategy z_a describes how player a chooses their move in each state. A strategy z_a for a player $a \in \Sigma$ induces a set of computations that the player a can enforce. Given a state $q \in Q$, a set $A \subseteq \Sigma$ of players, and a set $Z_A = \{z_a\}_{a \in A}$ of strategies, one strategy for each player in A , we define the outcomes of the strategies Z_A from q , denoted $out(q, Z_A)$, to be the set of q -computations that the players of A enforce following the strategies of Z_A . That is, a computation $\lambda = q_0, q_1, q_2, \dots$ is in $out(q, Z_A)$ if $q_0 = q$ and for all positions $i \geq 0$, there is a move vector $v = \langle j_1, \dots, j_k \rangle \in D$ such that

1. $j_a = z_a(q_i)$ for all players $a \in A$, and
2. $\delta(q_i, v) = q_{i+1}$.

The set $\mathcal{Z}^A = \{\{z_a\}_{a \in A} \mid z_a \text{ is a strategy for the player } a\}$ contains all sets of strategies, that contain one strategy for each player in $A \subseteq \Sigma$.

Alternating-time Temporal Logic The *alternating-time temporal logic* (ATL) [1] is defined with respect to a finite set Π of *propositions* and a finite set $\Sigma = \{1, \dots, k\}$ of players. An ATL formula is given by the abstract syntax:

$$\phi ::= p \mid \neg\phi \mid \phi_1 \vee \phi_2 \mid \langle\langle A \rangle\rangle\phi \mid \langle\langle A \rangle\rangle(\phi_1 \mathcal{U} \phi_2) \mid \llbracket A \rrbracket(\phi_1 \mathcal{U} \phi_2)$$

where $p \in \Pi$ is a proposition and $A \subseteq \Sigma$ is a set of players. The operators $\langle\langle \cdot \rangle\rangle$ and $\llbracket \cdot \rrbracket$ are *path quantifiers*. We will refer to them as "enforce" and "despite", respectively. The \bigcirc ("next") and \mathcal{U} ("until") are *temporal operators*. Additional temporal operators like \diamond ("eventually") and \square ("invariant") are derived as usual.

Given a state q of a game structure S , we write $q \models \phi$ to indicate that the state q satisfies the property described by ϕ . The satisfactory relation \models is defined inductively:

- $q \models p$ iff $p \in \pi(q)$.
- $q \models \neg\phi$ iff $q \not\models \phi$.
- $q \models \phi_1 \vee \phi_2$ iff $q \models \phi_1$ or $q \models \phi_2$.
- $q \models \langle\langle A \rangle\rangle\phi$ iff there exists a set $Z_A \in \mathcal{Z}^A$ of strategies, such that for all computations $\lambda \in out(q, Z_A)$, we have $\lambda[1] \models \phi$.
- $q \models \langle\langle A \rangle\rangle(\phi_1 \mathcal{U} \phi_2)$ iff there exists a set $Z_A \in \mathcal{Z}^A$ of strategies, such that for all computation $\lambda \in out(q, Z_A)$ there exists a position $i \geq 0$, such that $\lambda[i] \models \phi_2$ and for all positions $0 \leq j < i$, we have $\lambda[j] \models \phi_1$.
- $q \models \llbracket A \rrbracket(\phi_1 \mathcal{U} \phi_2)$ iff for all sets $Z_A \in \mathcal{Z}^A$ of strategies, we have that there exists a computation $\lambda \in out(q, Z_A)$, such that there exists a position $i \geq 0$, such that $\lambda[i] \models \phi_2$ and for all positions $0 \leq j < i$, we have $\lambda[j] \models \phi_1$.

3 Model Checking

In order to check if a CGS satisfies an ATL property, CGAAL encodes the problem as an extended dependency graph and finds a fixed-point assignment describing the satisfaction relation.

Extended Dependency Graphs An *extended dependency graph* (EDG) is a tuple $G = \langle C, E, N \rangle$ where C is a finite set of configurations (vertices), $E \subseteq C \times \mathcal{P}(C)$ is a finite set of hyper-edges, and $N \subseteq C \times C$ is a finite set of negation edges.

For a hyper-edge $e = \langle c, T \rangle \in E$ we call c the source configuration and $T \subseteq C$ the set of target configurations. Similarly, c is called the source configuration of the negation edge $\langle c, c' \rangle \in N$. We write $c \rightarrow c'$ if there exists an edge $\langle c, T \rangle \in E$ such that $c' \in T$ and $c \dashrightarrow c'$ if $\langle c, c' \rangle \in N$. Furthermore, we write $c \rightsquigarrow c'$ if $c \rightarrow c'$ or $c \dashrightarrow c'$. An EDG $G = \langle C, E, N \rangle$ is negation safe if there are no $c, c' \in C$ such that $c \dashrightarrow c'$ and $c' \rightsquigarrow^* c$. In what follows, we consider only negation-safe EDGs. Let $dist : C \rightarrow \mathbb{N}_0$ be the maximum number of negation edges throughout all paths starting in a configuration $c \in C$, inductively defined as:

$$dist(c) = \max\{dist(c'') + 1 \mid c', c'' \in C \text{ and } c \rightarrow^* c' \dashrightarrow c''\} \quad (1)$$

By convention $\max \emptyset = 0$. We define the $dist(G)$ of an EDG G as $dist(G) = \max_{c \in C} (dist(c))$. A component K_i of an EDG G , where $i \in \mathbb{N}_0$, is a subgraph induced on G by the sets

- $C_i = \{c \in C \mid dist(c) \leq i\}$
- $E_i = \{\langle c, T \rangle \in E \mid dist(c) \leq i\}$
- $N_i = \{\langle c, c' \rangle \in N \mid dist(c) \leq i\}$

denoting the set of configurations, hyper-edges and negation edges respectively in each respective component. By definition, the component K_0 has no negation edges.

Assignments An assignment $\alpha : C \rightarrow \{0, 1\}$ is a function that assigns boolean values to configurations of an EDG $G = \langle C, E, N \rangle$, where 0 and 1 represent false and true, respectively. We define α_0 as the zero assignment where $\alpha_0(c) = 0$ for all $c \in C$. We assume a component-wise ordering \sqsubseteq on assignments such that we have $\alpha \sqsubseteq \alpha'$ whenever $\alpha(c) \leq \alpha'(c)$ for all $c \in C$. The set of all assignments of an EDG G is denoted by \mathcal{A}^G and $(\mathcal{A}^G, \sqsubseteq)$ is a complete lattice.

By Knaster and Tarski's theorem, we can find a minimum fixed point on the complete lattice for any monotonic function. The minimum fixed-point assignment α_{\min} of an EDG G , denoted as $\alpha_{\min}^G = \alpha_{\min}^{K_{dist(G)}}$ is defined inductively on the components $K_0, K_1, \dots, K_{dist(G)}$ of G . For all $0 \leq i \leq dist(G)$, we define $\alpha_{\min}^{K_i}$ to be the minimum fixed-point assignment of the monotonic function $F_i : \mathcal{A}^{C_i} \rightarrow \mathcal{A}^{C_i}$ where

$$F_i(\alpha)(c) = \alpha(c) \vee \left[\bigvee_{\langle c, T \rangle \in E_i} \bigwedge_{c' \in T} \alpha(c') \right] \vee \left[\bigvee_{\langle c, c' \rangle \in N_i} \neg \alpha_{\min}^{K_{i-1}}(c') \right] \quad (2)$$

By convention, the conjunction of \emptyset is true and the disjunction of \emptyset is false. In the component K_0 there are no negation edges, which means the last clause in the disjunction is false for K_0 . We will refer to the repeated use of the F_i as the *global algorithm* because it requires building and iterating over the entire EDG. By Equation (2) each configuration in the EDG is comparable to a boolean formula in disjunctive normal form where each clause is another configuration or a negation of another configuration. Clearly, this makes EDGs an expressive structure and we shall now encode ATL model checking in an EDG.

3.1 Encoding of ATL in an EDG

We first establish some definitions related to subsets of move vectors induced by a coalition of players.

Partial Moves. Given a CGS S with the set of states Q and k players, we use $V_1 \times \dots \times V_k = \mathcal{V} \subseteq D(q)$ where $q \in Q$ to denote a *partial* move in q , where zero or more of the players' moves are fixed, i.e. their set of possible moves contain a single number. For instance, if $\mathcal{V} = \{2\} \times \{1,2,3\}$ in a 2-player game, then player 1 has chosen move 2 while player 2 is still free to choose from 1, 2, or 3. Furthermore, we use the notation $\mathcal{V}[a \mapsto j]$ for a variant \mathcal{V}' of a partial move $\mathcal{V} = V_1 \times \dots \times V_k$ where player a chooses move j . More precisely

$$\mathcal{V}[a \mapsto j] = V'_1 \times \dots \times V'_k \quad \text{where } 1 \leq i \leq k \text{ and } V'_i = \begin{cases} \{j\} & \text{if } i = a \\ V_i & \text{otherwise} \end{cases} \quad (3)$$

With this notation in mind, we define the function *pmoves*. The function *pmoves* gives the set of all possible partial moves that follow from the players of the set $A \subseteq \Sigma$ making a combination of moves in the state $q \in Q$:

$$\text{pmoves}(q, A) = \{D(q)[a \mapsto j_a][b \mapsto j_b] \dots \mid \{a, b, \dots\} = A \text{ and } \forall i \in A. 1 \leq j_i \leq d_i(q)\} \quad (4)$$

For instance, if $D(q) = \{1,2\} \times \{1,2\}$ in a 2-player game and $A = \{1\}$. Then $\text{pmoves}(q, A) = \{\{1\} \times \{1,2\}, \{2\} \times \{1,2\}\}$ contains two partial moves. One where player 1 chooses action 1, and another where player 1 chooses action 2. In other words, the *pmoves* function constructs the partial moves that the set A of players can choose from when *working together* in the state q . If A chooses $\mathcal{V} \in \text{pmoves}(q, A)$, then \mathcal{V} contains all move vectors that result from the remaining players $\Sigma \setminus A$ also making a choice.

We define a *partial transition function* Δ that produces a set of possible successor states given a state $q \in Q$ and a partial move \mathcal{V} . That is

$$\Delta(q, \mathcal{V}) = \{s \mid v \in \mathcal{V} \text{ and } \delta(q, v) = s\} \quad \text{where } \mathcal{V} \subseteq D(q) \quad (5)$$

With EDGs, assignments, and partial moves introduced, we can now define how the satisfaction relation \models can be encoded as an EDG.

Encoding. Given a CGS S , a state q of S , and an ATL state formula ϕ we now construct an EDG where every configuration is either a pair of a state and a formula, or a triple of a state, partial move, and a formula. The triples represent partially evaluated states, where some players' moves are already set. Starting from the initial pair $\langle q, \phi \rangle$, the dependency graph is constructed according to Figures 1 to 6. The figures show which outgoing edges each configuration has and the target configurations of those edges.

Theorem 3.1. *Given a CGS $S = \langle k, Q, \Pi, \pi, d, \delta \rangle$, a state $q \in Q$, and an ATL formula ϕ , let G be the negation-safe EDG rooted in $\langle q, \phi \rangle$ constructed following Figures 1 to 6. Then $\alpha_{\min}^G(\langle q, \phi \rangle) = 1$ if and only if $q \models \phi$.*

The proof of Theorem 3.1 can be found in the extended version of this paper.

Certain Zero In addition to the global algorithm (repeated application of F_i from Equation (2)), CGAAL also implements a local algorithm heavily inspired by the distributed CERTAINZERO algorithm by Dalsgaard et al. [4]. A local algorithm specialises in finding the assignment $\alpha_{\min}^G(\langle q, \phi \rangle)$ of a specific configuration $\langle q, \phi \rangle$ and not necessarily the assignment of the whole EDG. This allows the algorithm to terminate early in many cases. For the CERTAINZERO algorithm, Dalsgaard et al. also introduce a




Figure 1: EDG encoding of true, negation, and atomic proposition




Figure 2: EDG encoding of disjunction




Figure 3: EDG encoding of next




Figure 4: EDG encoding of enforce until




Figure 5: EDG encoding of despite until




Figure 6: EDG encoding of partially moved despite until



Figure 7: Ordering of assignment values for fixed-point computation in the CERTAINZERO algorithm. The top values are more certain.

symbolic assignment with new assignment values. The value \perp indicates that the configuration is unexplored. The value $?$ (unknown) indicates that the configuration is explored, but its final value has yet to be determined. Lastly, 0 and 1 indicate the final values in the minimum fixed-point assignment α_{\min}^G . This induces an ordering as seen in the lattice in Figure 7. During the algorithm, the assignments of the configurations will rise to more and more certain values, and if the assignment of a configuration ever becomes 0 or 1, we can be certain that its final value is 0 or 1, respectively. This is one of the properties that lead to improved performance, since in contrast to the global algorithm where all configurations are initially assigned 0, we can distinguish an initial 0 and a certain 0. Hence, the CERTAINZERO algorithm can terminate early if the root configuration is ever assigned either 0 or 1. Otherwise, it terminates when it can no longer raise any assignment values.

3.2 Search Strategies

The CERTAINZERO algorithm is controlled by a search strategy, which determines in which order the edges are explored and evaluated. Our search strategies for CGAAL include the common breadth-first search (BFS) and depth-first search (DFS) strategies, as well as multiple search strategies based on heuristics. Some of these are discussed in the following subsections. Which strategy is best depends heavily on the shape of the EDG which is determined by the CGS and the ATL formula in question. An evaluation of the strategies can be found in Section 5. The BFS strategy is the default strategy.

Dependency Heuristic Search (DHS) PageRank [15] is an algorithm that was created to estimate the importance of a website based on how many other websites have links to it. The idea has since been used in other areas, such as graph recommendation systems [13] or measuring structural-context similarity with SimRank [7]. Our dependency heuristic search (DHS) uses a similar idea by assuming that configurations with many ingoing edges are important. Finding the assignments of these, results in more back-propagation, bringing us closer to termination. In other words, the heuristic focuses on the *trunk* of the EDG where certain assignments are of high value. Specifically, DHS prioritise edges where the source configuration has a high number of ingoing edges. That is, if e is an edge with source configuration c then:

$$\text{priority}(e) = \text{indegree}(\text{source}(e)) = |\{e' \mid (e' = \langle c', T \rangle \in E \wedge c \in T) \vee e' = \langle c', c \rangle \in N\}| \quad (6)$$

Edges with the same priority are subject to FIFO ordering. However, since we do not know the whole graph in advance, we must continuously update the priority of edges when we explore successors of new configurations. This is only a small overhead with a priority queue data structure.

Linear Programming Heuristic Search (LPS) In this search strategy, we take advantage of how LCGS states can be represented as vectors, i.e. Q is a vector space. Given an edge with source configuration $\langle q, \phi \rangle$, we transform ϕ into a set \mathcal{L}_ϕ of linear constraints, each defined as a pair $\langle \mathbf{C}, b \rangle$, where \mathbf{C}




Figure 8: The distance $dist_{LPS}(\langle q, \phi \rangle) = 7$ when $q = \langle 7, 1 \rangle \in Q = \mathbb{R}^2$ and \mathcal{L}_ϕ corresponds to the constraints $-x + y \geq 1$ and $\frac{1}{2}x + y \geq 7$. The state s minimises $\|s - q\|_1$ while adhering to the constraints.

is a matrix and b is a vector. We prioritise edge e if its source configuration $\langle q, \phi \rangle$ has a low estimated distance to satisfaction as given by:

$$dist_{LPS}(\langle q, \phi \rangle) = \min \|s - q\|_1 \quad \text{subject to } \mathbf{C}s \geq b \quad \text{where } s \in Q, \langle \mathbf{C}, b \rangle \in \mathcal{L}_\phi \quad (7)$$

where $\|\cdot\|_1$ is the 1-norm, i.e. taxicab distance. The following equation is an equivalent linear programming problem:

$$\min \sum_i x_i \quad \text{subject to } \mathbf{C}s \geq b \text{ and } -x_i \leq s_i - q_i \leq x_i \quad \forall i \quad \text{where } s \in Q, \langle \mathbf{C}, b \rangle \in \mathcal{L}_\phi \quad (8)$$

A visualisation of the distance $dist_{LPS}$ can be seen on Figure 8, where $Q = \mathbb{R}^2$, $q = \langle 7, 1 \rangle \in Q$ and

$$\mathcal{L}_\phi = \left\{ \left\langle \begin{bmatrix} -1 & 1 \\ \frac{1}{2} & 1 \end{bmatrix}, \begin{bmatrix} 1 \\ 7 \end{bmatrix} \right\rangle \right\}$$

Many methods exist for solving linear programming problems. We use a library called `minilp`¹. The LPS strategy assumes value-wise close states are structurally close and works best when the constituent variables of states represent non-categorical data. The drawback of this method is the computational overhead of linear programming amplified by the potential need to solve multiple linear programming problems for each configuration based on the structure of ϕ . We reduce some of this overhead by caching the set \mathcal{L}_ϕ for each ϕ .

We also support an alternative search strategy involving linear programming. It is called Linear Representative Search (LRS) and it computes the distance described above for the root configuration only. The closest satisfying state s is then saved and edges are prioritised based on the 1-norm distance between s and the state in their source configuration. In other words, we assume that the state s found for the root represents all satisfied states. As a result this search strategy is cheaper than LPS but risks being inaccurate.

Instability Heuristic Search (IHS) The Petri Net model checker TAPAAL [8] implements the CERTAINZERO algorithm as well. Their configurations consist of a marking (a state) and a property, and their default search strategy uses a heuristic that estimates the distance between the marking of the configuration and a marking that satisfies the formula of the configuration. Our novel instability heuristic is inspired by their heuristic, but we differ by acknowledging that since there are negation edges in the EDG, we may not always be looking for a state that satisfies the formula. That is, if the state already

¹minilp crate: <https://crates.io/crates/minilp>

satisfies the associated formula, we estimate the distance to a state that does not instead. This guides the search towards configurations where the assignment is *unstable* and thus may have a high influence. Algorithms 1 and 2 find the $dist_{IHS}$ described above for an edge. The algorithms use an abstract metric $BiDist$ which for state-property pairs finds $\langle \hat{t}, \hat{f} \rangle$ where \hat{t} is the distance to the property being true and \hat{f} is the distance to the property being false.

Algorithm 1: $dist_{IHS}$

Input : An edge $e \in E$

Output: An estimated distance to a different assignment

```

1 if  $e = \langle c, T \rangle \in E$  then
2    $\langle \hat{t}, \hat{f} \rangle := \prod_{\langle q, \phi \rangle \in T} BiDist(q, \phi)$ ;
3   return  $\hat{t}$  if  $\hat{t} > 0$ , otherwise  $\hat{f}$ 
4 else if  $e = \langle c, \langle q, \phi \rangle \rangle \in N$  then
5    $\langle \hat{t}, \hat{f} \rangle := BiDist(q, \phi)$ ;
6   return  $\hat{f}$  if  $\hat{f} < 0$ , otherwise  $\hat{t}$ 

```

Algorithm 2: $BiDist$

Input : A state q and formula ϕ

Output: A bi-distance $\langle \hat{t}, \hat{f} \rangle$ describing the instability of ϕ in q

```

1 if  $\phi = expr_1 \triangleleft expr_2$  and  $v = eval(q, expr_1) - eval(q, expr_2)$  then
2   return  $\langle v, 0 \rangle$  if  $v > 0$ , otherwise  $\langle 0, v \rangle$ 
3 else if  $\phi = expr_1 \triangleright expr_2$  and  $v = eval(q, expr_2) - eval(q, expr_1)$  then
4   return  $\langle v, 0 \rangle$  if  $v > 0$ , otherwise  $\langle 0, v \rangle$ 
5 else if  $\phi = \neg \phi'$  and  $\langle \hat{t}, \hat{f} \rangle = BiDist(q, \phi')$  then
6   return  $\langle \hat{f}, \hat{t} \rangle$ 
7 else if  $\phi = \phi_1 \wedge \phi_2$  then
8   return  $BiDist(q, \phi_1) \sqcap BiDist(q, \phi_2)$ 
9 else if  $\phi = \phi_1 \vee \phi_2$  then
10  return  $BiDist(q, \phi_1) \sqcup BiDist(q, \phi_2)$ 
11 else if  $\phi = \langle\langle A \rangle\rangle \circ \phi'$  then
12  return  $BiDist(q, \phi')$ 
13 else if  $\phi = \langle\langle A \rangle\rangle (\phi_1 \mathcal{U} \phi_2)$  then
14  return  $BiDist(q, \phi_1) \sqcup BiDist(q, \phi_2)$ 
15 else if  $\phi = \llbracket A \rrbracket (\phi_1 \mathcal{U} \phi_2)$  then
16  return  $BiDist(q, \phi_1) \sqcup BiDist(q, \phi_2)$ 
17 // where  $\triangleleft \in \{<, \leq\}$ ,  $\langle \hat{t}_1, \hat{f}_1 \rangle \sqcap \langle \hat{t}_2, \hat{f}_2 \rangle = \langle \hat{t}_1 + \hat{t}_2, \min\{\hat{f}_1, \hat{f}_2\} \rangle$ 
18 //  $\triangleright \in \{>, \geq\}$ ,  $\langle \hat{t}_1, \hat{f}_1 \rangle \sqcup \langle \hat{t}_2, \hat{f}_2 \rangle = \langle \min\{\hat{t}_1, \hat{t}_2\}, \hat{f}_1 + \hat{f}_2 \rangle$ 

```

4 Tool Overview

CGAAL (<https://github.com/d702e20/CGAAL>) is written in Rust and consists of a command-line interface and a verification engine. The primary feature of CGAAL is the verification of ATL properties

for CGSs. Verification is either done with a global algorithm or local algorithm, both of which are described in Section 3. If requested, a partial strategy witness can also be computed, which instructs how the given players must play to satisfy the given property. Another feature converts the model into dot graph format such that it can be visually rendered with Graphviz.²

To model concurrent game structures for CGAAL, we designed a declarative Language for Concurrent Game Structure, called LCGS. In LCGS a CGS is defined with a series of declarations such that it is possible to find the successors of any state from the declarations alone. The language acts as an abstract representation of the CGS and allows us to save memory at the small cost of having to evaluate the expressions of declarations whenever successors are explored. The syntax of the language is inspired by PRISM-lang [11] used by the PRISM model checker to model stochastic multi-player games with rewards. However, LCGS differs in multiple ways. For instance, LCGS has player templates instead of modules, and templates have no effect unless there exists an instance of it. Additionally, synchronisations affecting the internal state of a player are much easier to declare.

Example Use As an example, we want to check if a cowboy can guarantee to stay alive in a three-way Mexican standoff. The standoff is simulated in rounds and in each round a cowboy can choose to wait, shoot the cowboy to the right, or shoot the cowboy to the left. If a cowboy is hit by two bullets, he dies.

```

1  const max_health = 2;
2
3  template cowboy
4
5      // How many bullets a cowboy can be hit by before dying
6      health : [0 .. max_health] init max_health;
7      health' = max(health - opp_right.shoot_left - opp_left.shoot_right, 0);
8
9      // A proposition used by ATL formulae
10     label alive = health > 0;
11
12     // The actions available to each cowboy
13     [wait] 1;
14     [shoot_right] health > 0 && opp_right.health > 0;
15     [shoot_left] health > 0 && opp_left.health > 0;
16
17 endtemplate
18
19 // The three cowboys in the Mexican standoff
20 player billy = cowboy [opp_right=clayton, opp_left=jesse];
21 player clayton = cowboy [opp_right=jesse, opp_left=billy];
22 player jesse = cowboy [opp_right=billy, opp_left=clayton];

```

Listing 1: LCGS implementation of a Mexican standoff

We model this scenario in Listing 1. The cowboy template is declared on lines 3-17. Herein, on line 6, we define that each cowboy can be hit by two bullets before being incapacitated, but we make this number a constant on line 1 so it is easy to change. Each cowboy has their own health variable and the combination of the values of these variables makes up a state. Line 7 defines how the health of a cowboy is updated in each transition. The value of `opp_right.shoot_left` and `opp_left.shoot_right` are 1 if this cowboy was shot by the given opponent to the right or left, respectively, otherwise 0. On line 10

²Graphviz: <https://graphviz.org/>

we define a label called `alive`, which is a proposition that is true if the cowboy has more than zero health. Lines 13-15 define the actions of the cowboy template, and the expression to the right of the name is a condition defining in which states the action is available. As can be seen, a cowboy can always wait, but only shoot if they and their target are alive. Lastly, we declare three instances of the cowboy template on lines 20-22. We define the right and left opponent in the relabelling function after the name of the template. Any identifier in a template can be relabelled to another expression as long as the result is syntactically and semantically correct.

The ATL formula $\langle\langle \text{billy} \rangle\rangle \Box \text{billy.alive}$ is satisfied if the cowboy `billy` has a strategy to stay alive. We can now run CGAAL with the following command: `./cgaal solver -m standoff.lcgs -f billy-can-stay-alive.atl`, and CGAAL will tell us that the property is not satisfied. Billy has no strategy that can guarantee his survival.

5 Evaluation

To evaluate our tool, we run several experiments. In our experiments we compare our global algorithm, our local algorithm using our various search strategies, and the established tool PRISM-games. We use several different concurrent game case studies. Some of these are PRISM-games case studies adapted to LCGS for CGAAL and determinised if needed. Others are well-known algorithms and games constructed during the development of CGAAL. We make the PRISM-lang and LCGS implementations of models as identical as possible, such that the state spaces are comparable. Some of the concurrent games and related ATL formulae used in the experiments are presented below. Queries marked with \top (resp. \perp) are satisfied (resp. not satisfied), and queries marked with \dagger may terminate early as we are not required to compute the entire fixed-point:

- **Mexican-Standoff:** In this model N cowboys stand in a circle, each with a revolver. At each moment they can choose to shoot another cowboy or do nothing. If a cowboy is hit by B bullets, they are incapacitated. We run the following queries:

$$\begin{aligned} MS_N^B \phi_1^{\perp\dagger} &= \langle\langle p_1 \rangle\rangle \Box p_{1_alive} \\ MS_N^B \phi_2^{\perp} &= \langle\langle p_1 \rangle\rangle \Diamond \neg p_{1_alive} \\ MS_N^B \phi_3^{\top} &= \langle\langle \{p_i \in \Sigma \mid 1 \equiv i \pmod{2}\} \rangle\rangle \Box p_{1_alive} \end{aligned}$$

- **Gossiping Girls:** In this model, N nodes each know a piece of information. The nodes can exchange information by calling each other. Our experiments include both a fully connected network of nodes and a circular network of nodes. We run the following queries:

$$\begin{aligned} GG_N \phi_1^{\top\dagger} &= \langle\langle \Sigma \rangle\rangle (calls \leq 10 \mathcal{U} all_girls_know_all) \\ GG_N \phi_2^{\top\dagger} &= \langle\langle \Sigma \rangle\rangle (calls \leq 10 \mathcal{U} only_p_1_knows_all) \\ GG_N \phi_3^{\top\dagger} &= \langle\langle \Sigma \rangle\rangle (calls \leq 10 \mathcal{U} p_1_knows_all) \\ GG_N \phi_4^{\top\dagger} &= \langle\langle \Sigma \rangle\rangle \Diamond calls > 10 \wedge all_girls_know_all \\ GG_N \phi_5^{\top\dagger} &= \langle\langle \emptyset \rangle\rangle \Diamond calls > 10 \\ GG_N \phi_6^{\perp\dagger} &= \langle\langle p_1 \rangle\rangle (calls \leq 10 \mathcal{U} p_1_knows_all) \\ GG_N \phi_7^{\top\dagger} &= \langle\langle \Sigma \setminus \{p_1\} \rangle\rangle \Diamond everyone_except_p_1_knows_all \end{aligned}$$

- **Robot Coordination:** In this model, four robots move orthogonally on a N by N grid. Each robot needs to reach the opposite corner of where they start, but crashing into each other is fatal. We run the following queries:

$$RC_N\phi_1^\perp = \langle\langle p_1, p_2, p_3 \rangle\rangle \diamond p_{1_at_target} \wedge p_{2_at_target}$$

$$RC_N\phi_2^\top = all_robots_at_home$$

$$RC_N\phi_3^{\top\dagger} = \langle\langle p_1, p_2, p_3, p_4 \rangle\rangle \diamond all_robots_at_targets$$

$$RC_N\phi_4^\perp = \langle\langle p_1, p_3, p_4 \rangle\rangle \diamond p_{1_at_target} \wedge \neg p_{1_crashed}$$

Each experiment is run with a time limit of two hours, allocated 128 GB of memory, and has 32 cores available regardless of how many worker threads CGAAL may spawn. All experiments are run on several identical AMD EPYC 7642 based servers, allowing only one experiment per node at a time to reduce noise in the results.


Results Here we provide a select set of experimental results that roughly exemplify our findings in general. For a more thorough set of results along with all of our experimental data and the Python script processing it, see the git repository <https://github.com/d702e20/evaluation-results/>.

We find that the local on-the-fly algorithm is often one order of magnitude faster than the global algorithm regardless of the search strategy. However, the global algorithm can compete with and sometimes outperforms the local algorithm in cases where the local algorithm cannot terminate early. This matches observations made by A. Dalsgaard et al. [4] and M. C. Jensen et al. [10].


PRISM-games is a stochastic model checker and checks queries involving both probability and rewards, capabilities which are irrelevant in our test cases. Therefore, in many cases, CGAAL's local algorithm is two orders of magnitude faster than PRISM-games, typically when the model has many synchronisations that affect internal states of modules such as in Mexican Standoff and Gossiping Girls as seen in Figures 9a and 9c. Such synchronisations also require a high amount of PRISM-lang code, while being easily expressed in LCGS. However, there are a few cases where PRISM-games is faster than CGAAL, e.g. $RC_N\phi_1^\perp$ and $RC_N\phi_4^\perp$ as seen in Figure 9b. Here the local algorithm cannot terminate early.

The choice of search strategy sometimes has a discernible impact on the performance of the local algorithm. Which search strategy is best varies from case to case and no definitive best strategy. The BFS strategy does not work well when the local algorithm can terminate early based on information found multiple moves into the model, but its low overhead is beneficial when early termination is not possible. The IHS strategy is often a good choice when early termination is possible. The LPS strategy is often significantly slower than the others, which is unsurprising given its overhead of solving linear programming problems. In models where differences in the states correspond to a distance in space, such as in robot coordination, the LPS strategy performs notably better and is even faster than both BFS and DFS on $RC_4\phi_3^{\top\dagger}$. The lightweight version, LRS, which only solves the linear problem once, matches LPS in these cases but is also generally better due to its lower overhead.


A general trend in our experiments is that we see an increase in the execution speed of our distributed implementation as we increase the number of compute threads available. The only times we do not see a speed-up with an increased number of compute threads are when the models are small enough such that the overhead of managing the additional threads is significant. Similarly, a single thread performs better than a few threads in many cases, since there is no communication overhead with only one compute thread. In general, depending on the search strategy employed, we see an improvement on a scale of one to two orders of magnitude when increasing the number of compute threads as exemplified in Figure 9d.




(a) Mexican-standoff: horizontal axis is (N, B) tuples with N being the number of cowboys in the model and B the number of bullets the cowboys can be hit by.



(b) Robot Coordination: horizontal axis is the size of the grid the robots are to manoeuvre on.



(c) Gossiping girls in a circular topology: horizontal axis being the number of girls in the circle.



(d) Gossiping girls with six girls in a circular topology: horizontal axis being the number of compute threads used in our distributed algorithm. Note that PRISM was always run with just one thread.

Figure 9: Select experimental results. CGAAL-results in Figures 9a to 9c show the best result for each search strategy when varying between 1 and 32 threads.

6 Conclusion

In this paper, we present CGAAL, our model checker of alternating-time temporal logic properties in concurrent games. CGAAL checks such properties by encoding the problem as an extended dependency graph and then computes the satisfaction relation using the distributed local on-the-fly CERTAINZERO algorithm by Dalsgaard et al. [4]. We provide multiple novel search strategies for the algorithm and allow concurrent games to be expressed in our language LCGS. Our experiments show that the local on-the-fly algorithm outperforms the global algorithm in the majority of cases. We also find that CGAAL outperforms the state-of-the-art tool PRISM-games by being up to two orders of magnitude faster, especially in models where synchronisations affect the internal state of modules and whenever we are not required to compute the entire fixed point. However, this comparison is unfair, since our test only uses a fraction of PRISM-games feature set.

CGAAL is still in early development and much work is needed before it competes with PRISM in terms of capabilities. However, dependency graphs have been used for encoding various model-checking problems, and we intend to incorporate these techniques into CGAAL.

References

- [1] Rajeev Alur, Thomas A. Henzinger & Orna Kupferman (2002): *Alternating-Time Temporal Logic*. *J. ACM* 49(5), p. 672–713, doi:10.1145/585265.585270.
- [2] Edmund M. Clarke & E. Allen Emerson (1982): *Design and synthesis of synchronization skeletons using branching time temporal logic*. In Dexter Kozen, editor: *Logics of Programs*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 52–71, doi:10.1007/BFb0025774.
- [3] Edmund M Clarke, Thomas A Henzinger, Helmut Veith, Roderick Bloem et al. (2018): *Handbook of model checking*. 10, Springer, doi:10.1007/978-3-319-10575-8.
- [4] Dalsgaard, Andreas E. and Enevoldsen, Søren and Fogh, Peter and Jensen, Lasse S. and Jepsen, Tobias S. and Kaufmann, Isabella and Larsen, Kim G. and Nielsen, Søren M. and Olesen, Mads Chr. and Pastva, Samuel and Srba, Jiří (2017): *Extended Dependency Graphs and Efficient Distributed Fixed-Point Computation*. In Wil van der Aalst & Eike Best, editors: *Application and Theory of Petri Nets and Concurrency*, Springer International Publishing, Cham, pp. 139–158, doi:10.1007/978-3-319-57861-3_10.
- [5] Søren Enevoldsen, Mathias Claus Jensen, Kim Guldstrand Larsen, Anders Mariegaard & Jiří Srba (2020): *Verification of Multiplayer Stochastic Games via Abstract Dependency Graphs*. *LOPSTR2020*, doi:10.1007/978-3-030-68446-4_13.
- [6] Søren Enevoldsen, Kim Guldstrand Larsen & Jiri Srba (2022): *Extended abstract dependency graphs*. *Int. J. Softw. Tools Technol. Transf.* 24(1), pp. 49–65, doi:10.1007/s10009-021-00638-8.
- [7] Glen Jeh & Jennifer Widom (2002): *SimRank: a measure of structural-context similarity*. In: *Proceedings of the Eighth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, July 23-26, 2002, Edmonton, Alberta, Canada*, ACM, pp. 538–543, doi:10.1145/775047.775126.
- [8] Jonas F. Jensen, Thomas Nielsen, Lars K. Oestergaard & Jiří Srba (2016): *TAPAAL and Reachability Analysis of P/T Nets*, pp. 307–318. Springer Berlin Heidelberg, Berlin, Heidelberg, doi:10.1007/978-3-662-53401-4_16.
- [9] Jonas Finnemann Jensen, Kim Guldstrand Larsen, Jiri Srba & Lars Kaerlund Østergaard (2016): *Efficient model-checking of weighted CTL with upper-bound constraints*. *International Journal on Software Tools for Technology Transfer* 18(4), pp. 409–426, doi:10.1007/s10009-014-0359-5.
- [10] Mathias Claus Jensen, Anders Mariegaard & Kim Guldstrand Larsen (2019): *Symbolic model checking of weighted PCTL using dependency graphs*. *NASA Formal Methods Symposium*, pp. 298–315, doi:10.1007/978-3-030-20652-9_20.

- [11] M. Kwiatkowska, G. Norman & D. Parker (2011): *PRISM 4.0: Verification of Probabilistic Real-time Systems*. In G. Gopalakrishnan & S. Qadeer, editors: *Proc. 23rd International Conference on Computer Aided Verification (CAV'11)*, LNCS 6806, Springer, pp. 585–591, doi:10.1007/978-3-642-22110-1_47.
- [12] M. Kwiatkowska, G. Norman, D. Parker & G. Santos (2020): *PRISM-games 3.0: Stochastic Game Verification with Concurrency, Equilibria and Time*. In: *Proc. 32nd International Conference on Computer Aided Verification (CAV'20)*, LNCS 12225, Springer, pp. 475–487, doi:10.1007/978-3-030-53291-8_25.
- [13] Sangkeun Lee, Minsuk Kahng & Sang goo Lee (2015): *Constructing compact and effective graphs for recommender systems via node and edge aggregations*. *Expert Systems with Applications* 42(7), pp. 3396–3409, doi:10.1016/j.eswa.2014.11.062.
- [14] Xinxin Liu & Scott A. Smolka (1998): *Simple Linear-Time Algorithms for Minimal Fixed Points (Extended Abstract)*. In Kim Guldstrand Larsen, Sven Skyum & Glynn Winskel, editors: *Automata, Languages and Programming, 25th International Colloquium, ICALP'98, Aalborg, Denmark, July 13-17, 1998, Proceedings, Lecture Notes in Computer Science* 1443, Springer, pp. 53–66, doi:10.1007/BFb0055040.
- [15] Lawrence Page, Sergey Brin, Rajeev Motwani & Terry Winograd (1998): *The PageRank Citation Ranking: Bringing Order to the Web*. Technical Report, Stanford Digital Library Technologies Project. Available at <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.31.1768>.