Extended Dependency Graphs and Efficient Distributed Fixed-Point Computation

Andreas E. Dalsgaard[1], Søren Enevoldsen[1], Peter Fogh[1], Lasse S. Jensen[1], Tobias S. Jepsen[1], Isabella Kaufmann[1], Kim G. Larsen[1], Søren M. Nielsen[1], Mads Chr. Olesen[1], Samuel Pastva[1,2], and Jiří Srba[1]

[1] Department of Computer Science, Aalborg University, Aalborg East, Denmark
[2] Faculty of Informatics, Masaryk University, Brno, Czech Republic

Abstract. Equivalence and model checking problems can be encoded into computing fixed points on dependency graphs. Dependency graphs represent causal dependencies among the nodes of the graph by means of hyper-edges. We suggest to extend the model of dependency graphs with so-called negation edges in order to increase their applicability. The graphs (as well as the verification problems) suffer from the state space explosion problem. To combat this issue, we design an on-the-fly algorithm for efficiently computing fixed points on extended dependency graphs. Our algorithm supplements previous approaches with the possibility to back-propagate, in certain scenarios, the domain value 0, in addition to the standard back-propagation of the value 1. Finally, we design a distributed version of the algorithm, implement it in an open-source tool, and demonstrate the efficiency of our general approach on the benchmark of Petri net models and CTL queries from the Model Checking Contest 2016.

1 Introduction

Model checking [9], a widely used verification technique for exhaustive state space search, may be both time and memory consuming as a result of the state space explosion problem. As a consequence, interesting real-life models can often be too large to be verified. Numerous approaches have been proposed to address this problem, including symbolic model checking and various abstraction techniques [7]. An alternative approach is to distribute the computation across multiple cores/machines, thus expanding the amount of available resources. Tools such as LTSmin [23] and DIVINE [1] have recently been exploring this possibility, without the need of being committed to a fixed model description language.

It has also been observed that model checking is closely related to the problem of evaluating fixed points [30, 20, 26, 6], as these are suitable for expressing system properties described in the logics like Computation Tree Logic (CTL) [8] or the modal μ-calculus [29]. This has been formally captured by the notion of dependency graphs of Liu and Smolka [30]. A dependency graph, consisting of a finite set of nodes and hyper-edges with multiple target nodes, is an abstract framework for efficient minimum fixed-point computation over the node

assignments that assign to each node the value 0 or 1. It has a variety of usages, including model checking [20, 26, 6] and equivalence checking [10]. Apart from formal verification, dependency graphs are also used to solve games based e.g. on timed game automata [5] or to encode Boolean equation systems [25].

Liu and Smolka proved in [30] that dependency graphs can be used to compute fixed points of Boolean graphs and to solve in linear time the P-complete problem HORNSAT [15]. They offered both a global and local algorithm for computing the minimum fixed-point value. The global algorithm computes the minimum fixed-point value for all nodes in the graph, though, we are often only interested in the values for some specific nodes. The advantage of the local algorithm is that it needs to compute the values only for a subset of the nodes in order to conclude about the assignment value for a given node of the graph. In practise, the local algorithm is superior to the global one [20] and to further boost its performance, we recently suggested a distributed implementation of the local algorithm with preliminary experimental results [10] conducted for weak bisimulation and simulation checking of CCS processes.

Our contributions. Neither the original paper by Liu and Smolka [30] nor the recent distributed implementation [10] handle the problem of negation in dependency graphs as this can break the monotonicity in the iterative evaluation of the fixed points. In our work, we extend dependency graphs with so-called *negation edges*, define a sufficient condition for the existence of unique fixed points and design an efficient algorithm for their computation, hence allowing us to encode richer properties rather than just plain equivalence checking or negation-free model checking. As we aim for a competitive implementation and applicability in various verification tools, it is necessary to offer the user the binary answer (whether a property holds or not or whether two systems are equivalent or not) together with the evidence why this is the case. This can be conveniently done by the use of *two-player games* between Attacker and Defender. In our implementation, it is possible for the user to play the role of Defender while the Attacker (played by the tool) can convince the user why a property does not hold. We formally define games played on the extended dependency graphs and prove a correspondence between the winner of the game and the fixed-point value of a node in a dependency graph.

In order to maximize the computation performance, we introduce a novel concept of *certain zero* value that can be back-propagated along hyper-edges and negation edges in order to ensure early termination of the fixed-point algorithm. This technique can often result in considerable improvements in the verification time and has not been, to the best of our knowledge, exploited in earlier work. To further enhance the performance, we present a *distributed algorithm* for a fixed-point computation and prove its correctness. Last but not least, we implement the distributed algorithm in an extensible open source framework and we demonstrate the applicability of the framework on CTL model checking of Petri nets. In order to do so, we integrate the framework into the tool TAPAAL [11, 21] and run a series of experiments on the Petri net models and queries from the Model Checking Contest (MCC) 2016 [27]. An early single-core

prototype of the tool implementing the negation edges and certain zero back-propagation also participated in the 2016 competition and was awarded a silver medal in the category of CTL cardinality verification with 23010 points, while the tool LoLa [32] (using a single core in the CTL category) took the gold medal with 27617 points. As documented by the experiments in this paper, our 4-core distributed algorithm now outperforms the optimized sequential algorithm and hence it will challenge the first place in the next year competition (also given that Lola employs stubborn set reduction techniques and query rewriting that were not yet used in our current implementation).

Related Work. Related algorithms for explicit distributed CTL model checking include the assumption based method [4] and a map-reduce based method [2]. Opposed to our algorithm, which computes a local result, these algorithms often focus on computing the global result. The local and global algorithms by Liu and Smolka [30] were also extended to weighted Kripke structures for weighted CTL model checking via symbolic dependency graphs [20], however, without any parallel or distributed implementation.

LTSmin [23] is a language independent model checker which provides a large amount of parallel and symbolic algorithms. To the best of our knowledge, LTSmin uses a symbolic algorithm based on binary decision diagrams for CTL model checking and even our sequential algorithm outperformed LTSmin at MCC'16 [27] (in e.g. CTL cardinality category LTSmin scored 12452 points compared to 23010 points achieved by our tool). Marcie [18] is another Petri net model checking tool that performs symbolic analysis using interval decision diagrams whereas our approach is based on explicit analysis using extended dependency graphs. Marcie was a previous winner of the CTL category at MCC'15 [28], however, in 2016 it finished on a third place with 18358 points after our tool and LoLa, the winner of the competition that we discussed earlier.

Other related work includes [3, 17, 22] designing parallel and/or distributed algorithms for model-checking of the alternation-free modal μ-calculus. As in our approach, they often employ the on-the-fly technique but our framework is more general as it relies on dependency graphs to which the various verification problems can be reduced. The notion of support sets as an evidence for the validity of CTL formulae has been introduced in [31] and it is close to a (relevant part of) assignment on a dependency graph, however, the game characterization of support sets was not further developed, as stated in [31]. In our work, we provide a natural game-theoretic characterization of an assignment on general dependency graphs and such a characterization can be used to provide an evidence about the fixed-point value of a node in a dependency graph.

Finally, there are several mature tools like FDR3 [14], CADP [13], SPIN [19] and mCRL2 [16], some of them implementing distributed and on-the-fly algorithms. The specification language of these is however often fixed and extensions of such a language requires nontrivial implementation effort. Our approach relies on reducing a variety of verification problems into extended dependency graphs and then on employing our optimized and efficient distributed implementation,

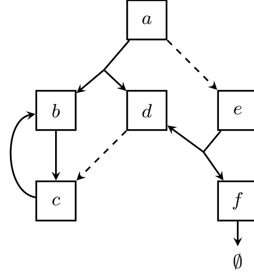

(a) An EDG with $dist(G) = 2$ and $V_0 = \{b, c, f\}$, $V_1 = \{d, e\} \cup V_0$, $V_2 = \{a\} \cup V_1$

	b	c	f
A_0	0	0	0
$F_0(A_0)$	0	0	1
$F_0(F_0(A_0))$	0	0	1

(b) $A^{C_0}_{min}$ Computation

	b	c	d	e	f
A_0	0	0	0	0	0
$F_1(A_0)$	0	0	1	0	1
$F_1(F_1(A_0))$	0	0	1	1	1
$F_1(F_1(F_1(A_0)))$	0	0	1	1	1

(c) $A^{C_1}_{min}$ Computation

	a	b	c	d	e	f
A_0	0	0	0	0	0	0
$F_2(A_0)$	0	0	0	1	1	1
$F_2(F_2(A_0))$	0	0	0	1	1	1

(d) $A^{C_2}_{min}$ Computation

Fig. 1: An EDG and iterative calculation of its minimum fixed-point assignment

as e.g. demonstrated on CTL model checking of Petri nets presented in this paper or on bisimulation checking of CCS processes [10].

2 Extended Dependency Graphs and Games

We shall now define the notion of extended dependency graphs, adding a new feature of negation edges to the original definition by Liu and Smolka [30].

Definition 1. *An Extended Dependency Graph (EDG) is a tuple $G = (V, E, N)$ where V is a finite set of* configurations, *$E \subseteq V \times \mathcal{P}(V)$ is a finite set of* hyper-edges, *and $N \subseteq V \times V$ is a finite set of* negation edges.

For a hyper-edge $e = (v, T) \in E$ we call v the *source configuration* and $T \subseteq V$ is the set of *target configurations*. We write $v \rightarrow u$ if there is a $(v, T) \in E$ such that $u \in T$ and $v \dashrightarrow u$ if $(v, u) \in N$. Furthermore, we write $v \rightsquigarrow u$ if $v \rightarrow u$ or $v \dashrightarrow u$. The *successor function* $succ : V \rightarrow (E \cup N)$ returns the set of outgoing edges from v, i.e. $succ(v) = \{(v, T) \in E\} \cup \{(v, u) \in N\}$. An example of an EDG is given in Figure 1(a) with the configurations named a to f, hyper-edges denoted by solid arrows with multiple targets, and dashed negation edges.

In what follows, we consider only EDGs without cycles containing negation edges.

Definition 2. *An EDG $G = (V, E, N)$ is* negation safe *if there are no $v, v' \in V$ s.t. $v \dashrightarrow v'$ and $v' \rightsquigarrow^* v$.*

After the restriction to negation safe EDG, we can now define the negation *distance function* $dist : V \rightarrow \mathbb{N}_0$ that returns the maximum number of negation edges throughout all paths starting in a configuration v and is inductively defined as $dist(v) = \max(\{dist(v'') + 1 \mid v', v'' \in V \text{ and } v \rightarrow^* v' \dashrightarrow v''\})$ where by convention $\max(\emptyset) = 0$. Note that $dist(v)$ is always finite because every path can visit each negation edge at most once. We then define $dist(G)$ of an EDG G as $dist(G) = max_{v \in V}(dist(v))$ and for an edge $e \in E \cup N$ where v is its source configuration, we write $dist(e) = dist(v)$.

A *component* C_i of G, where $i \in \mathbb{N}_0$, is a subgraph induced on G by the set of configurations $V_i = \{v \in V \mid dist(v) \leq i\}$. We write V_i, E_i and N_i to denote the set of configurations, hyperedges and negation edges of each respective component. Note that by definition, C_0 does not contain any negation edges. Also observe that $G = C_{dist(G)}$ and that for all $k, \ell \in \mathbb{N}_0$, if $k < \ell$ then C_k is a subgraph of C_ℓ. The EDG G in our example from Figure 1(a) contains three nonempty components and has $dist(G) = 2$.

An *assignment* A of an EDG $G = (V, E, N)$ is a function $A : V \rightarrow \{0, 1\}$ that assigns the value 0 (interpreted as false) or the value 1 (interpreted as true) to each configuration of G. A *zero assignment* A_0 is such that $A_0(v) = 0$ for all $v \in V$. We also assume a component wise ordering $\sqsubseteq$ of assignments such that $A_1 \sqsubseteq A_2$ whenever $A_1(v) \leq A_2(v)$ for all $v \in V$. The set of all assignments of G is denoted by $\mathcal{A}^G$ and clearly $(\mathcal{A}^G, \sqsubseteq)$ is a complete lattice.

We are now ready to define the minimum fixed-points assignment of an EDG G (assuming that a conjunction over the empty set is true, while a disjunction over the empty set is false).

Definition 3. *The* minimum fixed-point assignment *of an EDG G, denoted by $A^G_{min} = A^{C_{dist(G)}}_{min}$ is defined inductively on the components $C_0, C_1, \ldots, C_{dist(G)}$ of G. For all i, $0 \leq i \leq dist(G)$, we define $A^{C_i}_{min}$ to be the minimum fixed-point assignment of the function $F_i : \mathcal{A}^{C_i} \rightarrow \mathcal{A}^{C_i}$ where*

$$F_i(A)(v) = A(v) \vee \left[\bigvee_{(v,T) \in E_i} \bigwedge_{u \in T} A(u) \right] \vee \left[\bigvee_{(v,u) \in N_i} \neg A^{C_{i-1}}_{min}(u) \right] . \qquad (1)$$

Note that when computing the minimum fixed-point assignment $A^{C_0}_{min}$ for the base component C_0, we know that $N_0 = \emptyset$ and hence the third disjunct in the function F_0 always evaluates to false. In the inductive steps, the assignment $A^{C_{i-1}}_{min}$ is then well defined for the use in the function F_i. It is also easy to observe that each function F_i is monotonic (by a simple induction on i) and hence by Knaster-Tarski, the unique minimum fixed-point always exists for each i.

In Figure 1 we show the iterative computation of $A^{C_0}_{min}$, $A^{C_1}_{min}$ and $A^{C_2}_{min}$, starting from the zero assignment A_0. We iteratively upgrade the assignment of a configuration v from the value 0 to 1 whenever there is a hyper-edge (v, T) such that all target configurations $u \in T$ already have the value 1 or whenever there is a negation edge $v \dashrightarrow u$ such that the minimum fixed-point assignment of u (computed earlier) is 0. Once the application of the function F_i stabilizes, we have reached the minimum fixed-point assignment for the component C_i.

Remark 1. The algorithm for computing $A^{C_i}_{min}$ described above, also called the *global algorithm*, relies on the fact that the complete minimum fixed-point assignment of smaller components C_j where $j < i$ must be available before we can proceed with the computation on the component C_i. As we show later on, it is not always necessary to know the whole $A^{C_{i-1}}_{min}$ in order to compute $A^{C_i}_{min}(v)$ for a specific configuration v and such a computation can be done in an on-the-fly manner, using the so-called *local algorithm*.

2.1 Game Characterization

In order to offer a more intuitive understanding of the minimum fixed-point computation on an extended dependency graph G, and to provide a convincing argumentation why the minimum fixed-point value in a given configuration v is 0 or 1 (for the use in our tool), we define a two player game between the players *Defender* and *Attacker*. The *positions* of the game are of the form (v, r) where $v \in V$ is a configuration and $r \in \{0, 1\}$ is a claim about the minimum fixed-point value in v, postulating that $A^G_{min}(v) = r$. The game is played in *rounds* and Defender defends the current claim while Attacker does the opposite.

Rules of the Game: In each round starting from the current position (v, r), the players determine the new position for the next round as follows:

- If $r = 1$ then Defender chooses an edge $e \in succ(v)$. If no such edge exists then Defender loses, otherwise
  - if $e = (v, u) \in N$ then $(u, 0)$ becomes the new current position, and
  - if $e = (v, T) \in E$ then Attacker chooses the next position $(u, 1)$ where $u \in T$, unless $T = \emptyset$ which means that Attacker loses.
- If $r = 0$ then Attacker chooses an edge $e \in succ(v)$. If no such edge exists then Attacker loses, otherwise
  - if $e = (v, u) \in N$ then $(u, 1)$ becomes the new current position, and
  - if $e = (v, T) \in E$ then Defender chooses the next position $(u, 0)$ where $u \in T$, unless $T = \emptyset$ which means that Defender loses.

A *play* is a sequence of positions formed according the rules of the game. Any finite play is lost either by Defender or Attacker as defined above. If a play is infinite, we observe that the claim r can be switched only finitely many times (since the graph is negation safe). Therefore there is only one claim r that is repeated infinitely often in such a play. If $r = 1$ is the infinitely repeated claim then Defender loses, otherwise ($r = 0$) Attacker loses.

The game starting from the position (v, r) is *winning for Defender* if she has a universal winning strategy from (v, r). Similarly, the position is *winning for Attacker* if he has a universal winning strategy from (v, r). Clearly, the game is determined such that only one of the players has a universal winning strategy and from the symmetry of the game rules, we can also notice that Defender is the winner from (v, r) if and only if Attacker is the winner from $(v, 1 - r)$.

Theorem 1. *Let G be a negation safe EDG, $v \in V$ be a configuration and $r \in \{0, 1\}$ be a claim. Then $A^G_{min}(v) = r$ if and only if Defender is the winner of the game starting from the position (v, r).*

Proof. By induction on the level of the node v (a node is of level i if it belongs to the component C_i but not to any component C_j where $j < i$), followed by a case analysis. □

Let us now argue that Defender wins from the position $(a, 0)$ in the EDG G from Figure 1(a). First, Attacker picks either (i) the hyper-edge $(a, \{b, d\})$ or (ii) the negation edge (a, e). In case (i), Defender answers by selecting the configuration b and the game continues from $(b, 0)$. Now Attacker can only pick the hyper-edge $(b, \{c\})$ and Defender is forced to select the configuration c, ending in the position $(c, 0)$ and from here the only continuation of the game brings us again to the position $(b, 0)$. As the play now repeats forever with the claim 0 appearing infinitely often, Defender wins this play. In case (ii) where Attacker selects the negation edge, we continue from the position $(e, 1)$. Defender is forced to select the only available hyper-edge $(e, \{d, f\})$, on which Attacker can answer by selecting the new position $(d, 1)$ or $(f, 1)$. The first choice is not good for Attacker, as Defender will answer by taking the negation edge (d, c) and ending in the position $(c, 0)$ from which we already know that Defender wins. The position $(f, 1)$ is not good for Attacker either as Defender can now select the hyper-edge $(f, \emptyset)$ and Attacker loses as he gets stuck. Hence Defender has a universal winning strategy from $(a, 0)$ and by Theorem 1 we get that $A^G_{min}(a) = 0$.

2.2 Encoding of CTL Model Checking of Petri Nets into EDG

We shall now give an example of how CTL model checking of Petri nets can be encoded into computing fixed-points on EDGs. Let us first recall the Petri net model. Let $\mathbb{N}_0$ denote the set of natural numbers including zero and $\mathbb{N}_\infty$ the set of natural numbers including infinity.

A *Petri net* is a 4-tuple $N = (P, T, F, I)$ where P is a finite set of places, T is a finite set of transitions such that $P \cap T = \emptyset$ and $P \cup T \neq \emptyset$, $F : (P \times T \cup T \times P) \to \mathbb{N}_0$ is the flow function and $I : P \times T \to \mathbb{N}_\infty$ is the inhibitor function. A *marking* on N is a function $M : P \to \mathbb{N}_0$ assigning a number of tokens to each place. The set of all markings on N is denoted $M(N)$. A transition t is enabled in a marking M if $M(p) \geq F((p, t))$ and $M(p) < I(p, t)$ for all $p \in P$. If t is enabled in M, it can fire and produce a marking M', written $M \xrightarrow{t} M'$, such that $M'(p) = M(p) - F((p, t)) + F((t, p))$ for all $p \in P$. We write $M \to M'$ if there is $t \in T$ such that $M \xrightarrow{t} M'$.

In CTL, properties are expressed using a combination of logical and temporal operators over a set of basic propositions. In our case the propositions express properties of a concrete marking M and include the proposition **is_fireable**(Y) for a set of transitions Y that is true iff at least one of the transitions from Y is enabled in the marking M, and arithmetical expressions and predicates over the basic construct **token_count**(X) where X is a subset of places such that **token_count**(X) returns the total number of tokens in the places from the set X in the marking M. The CTL logic is motivated by the requirements of the MCC'16 competition [27] and the syntax of CTL formula φ is

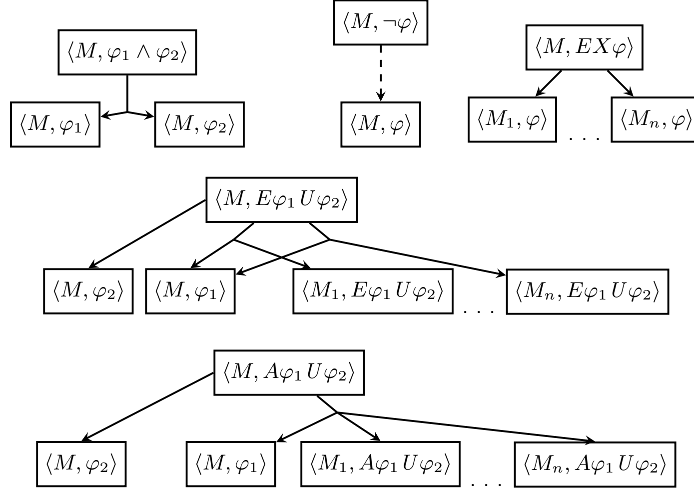

Fig. 2: Construction of EDG where we let $\{M_1, ..., M_n\} = \{M' \mid M \to M'\}$

$$\varphi ::= \mathit{true} \mid \mathit{false} \mid \textbf{is_fireable}(Y) \mid \psi_1 \bowtie \psi_2 \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid$$
$$EG\ \varphi \mid AG\ \varphi \mid EF\ \varphi \mid AF\ \varphi \mid EX\ \varphi \mid AX\ \varphi \mid E\varphi_1 U \varphi_2 \mid A\varphi_1 U \varphi_2$$
$$\psi ::= \psi_1 \oplus \psi_2 \mid c \mid \textbf{token_count}(X)$$

where $\bowtie \in \{<, \leq, =, \geq, >\}$, $X \subseteq P$, $Y \subseteq T$, $c \in \mathbb{N}_0$ and $\oplus \in \{+, -, \cdot\}$. We assume the standard semantics of satisfability of a CTL formula φ in a marking M, written $M \models \varphi$.

For a given marking M and a CTL formula φ, we now construct an EDG with the configurations of the form $\langle M, \varphi\rangle$. If φ is a basic proposition then there is a hyper-edge from $\langle M, \varphi\rangle$ with the empty target set iff $M \models \varphi$. The rules for building EDG for a subset of the temporal operators (the other temporal operators are the derived ones) is given in Figure 2. Observe that this reduction produces a negation safe EDG. We can then conclude with the following correctness result.

Theorem 2 (Encoding Correctness). *Let N be a Petri net, M a marking on N and let φ be a CTL formula. Let G be the EDG with the root $\langle M, \varphi\rangle$ constructed as described above. Then $M \models \varphi$ iff $A^G_{min}(\langle M, \varphi\rangle) = 1$.*

Remark 2. The reader probably noticed that if the Petri net is unbounded (has infinitely many reachable markings), we are actually producing an infinite EDG. Indeed, CTL model checking for unbounded Petri nets is undecidable [12], so we cannot hope for a general algorithmic solution. However, due to the employment of our local algorithm with certain zero propagation, we are sometimes able to obtain a conclusive answer by exploring only a finite part of the (on-the-fly) constructed dependency graph.

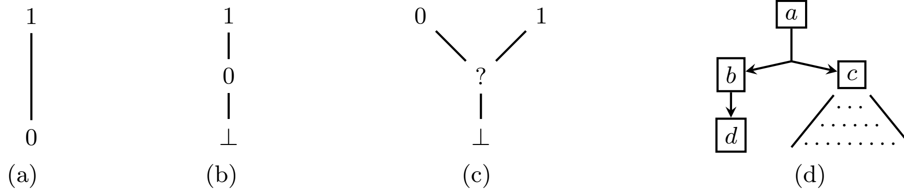

Fig. 3: Comparison of Different Algorithms for Fixed-Point Computation

3 Algorithms for Fixed-Point Computation on EDG

We shall now discuss the differences of our new distributed algorithm for fixed-point computation of EDG compared to the previous approaches, followed by the description of our algorithm.

Figure 3 shows the partial ordering of the assignment values used by the algorithms. The orderings in the figure show how the configuration values are upgraded during the execution of the algorithms. The global algorithm, described in Section 2, only uses the assignment values 0 and 1 as shown in Figure 3(a). Initially, the whole graph is constructed and all configurations are assigned the value 0. Then it iterates, starting from the component C_0, over all hyper-edges and upgrades the source configuration values to 1 whenever all target configurations are already assigned the value 1. This repeats until no further upgrades are possible and then it uses the negation edges to propagate the values to the higher components until the minimum fixed-point assignment of a given configuration is set to 1 (in which case an early termination is possible) or until the whole process terminates and we can claim that the minimum fixed-point assignment of the given configuration is 0.

The key insight for the local algorithm, as suggested by Liu and Smolka [30] for dependency graphs without negation edges, is that if we are only interested in $A^G_{min}(v)$ for a given configuration v, we do not have to necessarily enumerate the whole graph and compute the value for all configurations in G in order establish that $A^G_{min}(v) = 1$. The local algorithm introduces the value $\bot$ for not yet explored configurations as shown in Figure 3(b) and performs a forward search in the dependency graph with backward propagation of the value 1. This significantly improves the performance of the global algorithm in case the configuration v gets the value 1. In the case where $A^G_{min}(v) = 0$, the local algorithm must search the whole graph before terminating and announcing the final answer.

Our improvement to the local algorithm is twofold: the handling of negation edges in an on-the-fly manner and the introduction of a new value ?, taking over the previous role of 0, as shown in Figure 3(c). Here $\bot$ means that a configuration has not been discovered yet, ? that the final minimum fixed-point assignment has not been determined yet, and 0 and 1 mean the final values in the minimum fixed-point assignment. Hence as soon as the given configuration gets the value 0 or 1,

we can early terminate and announce the answer. The previous approaches did not allow early termination for the value 0, but as Figure 3(d) shows, it can save lots of work. Since d has no outgoing hyper-edges, it can get assigned the value 0 (called *certain zero*) and because the single target configuration of the hyper-edge $(b, \{d\})$ is 0, the value 0 can back-propagate to b (we do this by removing hyper-edges that contain at least one target configuration with the value 0 and once a configuration has no outgoing hyper-edges, it will get assigned the certain zero value 0). Now the hyper-edge $(a, \{b, c\})$ can also be removed and as a no longer has any hyper-edges, we can conclude that $A^G_{min}(a) = 0$ without having to explore the potentially large subgraph rooted at c as it would be necessary in the previous algorithms. We moreover have to deal with negation edges where we allow early back-propagation of the certain 0 and certain 1 values, essentially performing an on-the-fly search for the existence of Defender's winning strategy. In what follows, we shall present the formal details of our algorithm, including its distributed implementation.

3.1 Distributed Algorithm for Minimum Fixed-Point Computation

We assume n workers running Algorithm 1 in parallel. Each worker has a unique identifier $i \in \{1, ..., n\}$ and can communicate with any other worker using order preserving, reliable channels. If not stated otherwise, i refers to the identifier of the local worker and j refers to an identifier of some remote worker.

Global Data Structures. Initially, each worker has access to the means of generating a given EDG $G = (V, E, N)$ via the function *succ*, an initial configuration $v_0 \in V$, and a partition function $\delta : V \to \{1, \ldots, n\}$ that splits the configurations among the workers. We say that worker i owns a configuration v if $\delta(v) = i$.

Local Data Structures. Each worker has the following local data structures:
- $W^i_E \subseteq E$ is the waiting list of hyper-edges,
- $W^i_N \subseteq N$ is the waiting list of negation edges,
- $D^i : V \to \mathcal{P}(E \cup N)$ is the dependency set for each configuration,
- $succ^i : V \to \mathcal{P}(E \cup N)$ is the local successor relation such that initially $succ^i(v) = succ(v)$ if $\delta(v) = i$ and otherwise $succ^i(v) = \emptyset$,
- $A^i : V \to \{\perp, ?, 0, 1\}$ is the assignment function (implemented via hashing), initially returning $\perp$ for all configurations,
- $C^i : V \to \mathcal{P}(\{1, \ldots, n\})$ is the set of interested workers who requested the value of a given configuration,
- $M^i_R \subseteq V \times \{1, \ldots, n\}$ is the (unordered) message queue for requests (v, j), where j is the identifier of the worker requesting the assigned value (i.e. 0 or 1) of a configuration v belonging to the partition of worker i, and
- $M^i_A \subseteq V \times \{0, 1\}$ is the (unordered) message queue for answers (v, a), where a is the assigned value of configuration v which has been previously requested by worker i.

For syntactical convenience, we assume that we can add messages to M^i_R and M^i_A directly from other workers.

Algorithm 1 Distributed Certain Zero Algorithm for a Worker i

Require: Worker id i, an EDG $G = (V, E, N)$ and an initial configuration $v_0 \in V$.
Ensure: The minimum fixed-point assignment $A^G_{min}(v_0)$

```
 1: function DistributedCertainZero(G, v0)
 2:     if δ(v0) = i then Explore(v0)                                   ▷ Algorithm 2
 3:     repeat
 4:         if W_E^i ∪ W_N^i ∪ M_R^i ∪ M_A^i ≠ ∅ then
 5:             task ← PickTask(W_E^i, W_N^i, M_R^i, M_A^i)
 6:             if task ∈ W_E^i then ProcessHyperEdge(task)              ▷ Algorithm 2
 7:             else if task ∈ W_N^i then ProcessNegationEdge(task)     ▷ Algorithm 2
 8:             else if task ∈ M_R^i then ProcessRequest(task)          ▷ Algorithm 2
 9:             else if task ∈ M_A^i then ProcessAnswer(task)           ▷ Algorithm 2
10:     until TerminationDetection
11:     if A^i(v0) = ? ∨ A^i(v0) = 0 then return 0
12:     else return 1
```

Global waiting lists. When we need to reference the global state in the computation of the parallel algorithm, we can use the following abbreviations.

- The global waiting list of hyper-edges $W_E = \bigcup_{i=1}^{n} W_E^i$.
- The global waiting list of negation edges $W_N = \bigcup_{i=1}^{n} W_N^i$.
- The global request message queue $M_R = \bigcup_{i=1}^{n} M_R^i$.
- The global answer message queue $M_A = \bigcup_{i=1}^{n} M_A^i$.

Idle Worker. We say that a worker i is idle if it is executing the loop at line 3 through 10 in Algorithm 1, but is not currently executing any of the processing functions on lines 6, 7, 8 or 9.

Pick Task. Algorithm 1 uses at line 5 the function PickTask($W_E^i, W_N^i, M_R^i, M_A^i$) that nondeterministically returns:

- a hyper-edge from W_E^i, or
- a message from M_R^i or M_A^i, or
- a negation edge (v, u) from W_N^i provided that $A^i(u) \in \{0, 1, \bot\}$, or
- a negation edge (v, u) from W_N^i if all other workers are idle and v has a minimal distance in all waiting lists and message queues (i.e. for all $(v', x) \in (W_E \cup W_N \cup M_A \cup M_R)$ it holds that $dist(v) \leq dist(v')$).
- If none of the above is satisfied, the worker waits until either a message is received or a negation edge becomes safe to pick. Notice that in this case, W_E^i will remain empty until a message or negation edge is processed.

Even though PickTask depends on the global state of the computation to decide whether a negation edge is safe to pick, the rest of the conditions can be decided based on the data that is available locally to each worker. Therefore it is not necessary to synchronise across all workers every time a task should be picked, it is only required if the worker wants to pick a negation edge (v, u) where $A^i(u) = ?$.

Termination of the Algorithm. We utilize a standard TerminationDetection function computed distributively that returns *true* if and only if all message

Algorithm 2 Functions for Worker i Called from Algorithm 1

```
1: function ProcessHyperEdge(e = (v, T))                                   ▷ e ∈ E
2:     W_E^i ← W_E^i \ {e}
3:     if ∀u ∈ T : A^i(u) = 1 then FinalAssign(v, 1)          ▷ Edge propagates 1
4:     else if ∃u ∈ T where A^i(u) = 0 then DeleteEdge(e)
5:     else if X ⊆ T s.t. X ≠ ∅ and ∀u ∈ X : A^i(u) = ? ∨ A^i(u) = ⊥ then
6:         for u ∈ X do
7:             D^i(u) ← D^i(u) ∪ {e}
8:             if A^i(u) = ⊥ then Explore(u)

1: function ProcessNegationEdge(e = (v, u))                                ▷ e ∈ N
2:     W_N^i ← W_N^i \ {e}
3:     if A^i(u) = ? ∨A^i(u) = 0 then FinalAssign(v, 1)      ▷ Assign negated value
4:     else if A^i(u) = 1 then DeleteEdge(e)
5:     else if A^i(u) =⊥ then
6:         D^i(u) ← D^i(u) ∪ {e}; W_N^i ← W_N^i ∪ {e}; Explore(u)

1: function ProcessRequest(m = (v, j))                     ▷ request from worker j
2:     if A^i(v) = 1 ∨ A^i(v) = 0 then               ▷ Value of v is already known
3:         M_A^j ← M_A^j ∪ {(v, A^i(v))} ; M_R^i ← M_R^i \ {m}
4:     else                                          ▷ Value of v is not computed yet
5:         C^i(v) ← C^i(v) ∪ {j}             ▷ Remember that worker j is interested in v
6:         M_R^i ← M_R^i \ {m}
7:         if A^i(v) = ⊥ then Explore(v)

1: function ProcessAnswer(m = (v, a))                  ▷ a ∈ {0, 1} and m ∈ M_A^i
2:     M_A^i ← M_A^i \ {m}
3:     FinalAssign(v, a)                          ▷ Assign the received answer to v

1: function Explore(v)                                                     ▷ v ∈ V
2:     A^i(v) ← ?
3:     if δ(v) = i then                                      ▷ Does worker i own v?
4:         if succ^i(v) = ∅ then FinalAssign(v, 0)         ▷ It is safe to propagate 0
5:         W_E^i ← W_E^i ∪ (succ^i(v) ∩ E); W_N^i ← W_N^i ∪ (succ^i(v) ∩ N)
6:     else
7:         M_R^δ(v) ← M_R^δ(v) ∪ {(v, i)}    ▷ If not, request the value from the owner of v

1: function DeleteEdge(e = (v, T) or e = (v, u))                     ▷ e ∈ (E ∪ N)
2:     succ^i(v) ← succ^i(v) \ {e}
3:     if succ^i(v) = ∅ then FinalAssign(v, 0)             ▷ It is safe to propagate 0
4:     if e ∈ E then
5:         W_E^i ← W_E^i \ {e}
6:         for all u ∈ T do D^i(u) ← D^i(u) \ {e}
7:     if e ∈ N then
8:         W_N^i ← W_N^i \ {e}; D^i(u) ← D^i(u) \ {e}

1: function FinalAssign(v, a)                           ▷ a ∈ {0, 1} and v ∈ V
2:     if v = v_0 then return a and terminate all workers;       ▷ Early termination
3:     A^i(v) ← a
4:     for all j ∈ C^i(v) do M_A^j ← M_A^j ∪ {(v, a)}    ▷ Notify all interested workers
5:     W_E^i ← W_E^i ∪ {D^i(v) ∩ E}; W_N^i ← W_N^i ∪ {D^i(v) ∩ N}
```

queues are empty, all waiting lists are empty (i.e. $W_E \cup W_N \cup M_R \cup M_A = \emptyset$) and all workers are idle. Notice that once the initial configuration v_0 is assigned the final value 0 or 1, the algorithm can terminate early.

We shall now focus on the correctness of the algorithm. By a simple code analysis, we can observe the following lemma.

Lemma 1. *During the execution of Algorithm 1, the value of $A^i(v)$ for any worker i and any configuration v will never decrease (with respect to the ordering from Figure 3(c)).*

Based on this lemma we can now argue about the termination of the algorithm.

Lemma 2. *Algorithm 1 terminates.*

Proof. To show that the algorithm terminates, we have to argue that eventually all waiting lists become empty and all workers go to idle (unless early termination kicks in before this). By guaranteeing this, the TerminationDetection condition will be satisfied and the algorithm terminates.

First, let us observe that if the waiting lists of a worker are empty, the worker will eventually become idle. That is because none of the functions called from the repeat-until loop contain any loops or recursive calls. Also note that in such case, the worker will stay idle until a message is received. In each iteration, an edge is inserted into a waiting list only if the assignment value of some configuration increases. By Lemma 1, the assignment value can never decrease, and since the assignment value can only increase finitely many times, eventually no edges will be inserted into the waiting lists. The same argument applies to request messages as a request can only be sent if an assignment value of a configuration increases from $\perp$ to ?. The only exception to the considerations above are the answer messages. An answer message can be sent either as a result of an assignment value increase (line 4 of the FinalAssign), which only happens finitely many times. However, it can be also sent as a direct response to a request message (line 3 of the ProcessRequest). As we have already shown, each computation can produce only finitely many requests and since each such request can produce at most one answer, the number of answer messages will also be finite.

Finally, we note that as soon as all the messages and hyper-edges are processed by all workers, at least one negation edge becomes safe to pick. Hence if no new messages are sent or edges being inserted into the waiting lists, eventually a negation edge is picked (at most once). Therefore all waiting lists become eventually empty and as a result all workers go idle, satisfying the TerminationDetection condition. □

The main correctness argument is contained in the following loop invariants.

Lemma 3 (Loop Invariants). *For any worker i, the repeat-until loop in Algorithm 1 satisfies the following invariants.*

1. *For all $v \in V$, if $A^i(v) = 1$ then $A^G_{min}(v) = 1$.*

2. *For all* $v \in V$, *if* $A^i(v) = 0$ *then* $A^G_{min}(v) = 0$.
3. *For all* $v \in V$, *if* $A^i(v) = ?$ *and* $i = \delta(v)$ *then for all* $e \in succ^i(v)$ *it holds that* $e \in W^i_E \cup W^i_N$ *or* $e \in D^i(u)$ *for some* $u \in V$ *where* $A^i(u) = ?$.
4. *For all* $v \in V$, *if* $A^i(v) = ?$ *and* $i \neq \delta(v)$ *then one of the following must hold:*
   - $(v, i) \in M^{\delta(v)}_R$,
   - $i \in C^{\delta(v)}(v)$ *and* $A^{\delta(v)}(v) = ?$, *or*
   - $(v, a) \in M^i_A$ *and* $A^{\delta(v)}(v) = a$ *for some* $a \in \{0, 1\}$.
5. *If there is a negation edge* $e = (v, u) \in W^i_N$ *s.t.* $A^i(u) = ?$ *and all workers are idle and* v *is minimal in all waiting lists and message queues (i.e. for all* $(v', x) \in (W_E \cup W_N \cup M_A \cup M_R)$ *it holds that* $dist(v) \leq dist(v')$*), then* $A^G_{min}(u) = 0$.

Now we can state two technical lemmas.

Lemma 4. *Upon termination of Algorithm 1 at line 11 or line 12, for every negation edge* $e = (v, u) \in N$ *it holds that either* $A^{\delta(v)}(v) \in \{1, \perp\}$ *or the negation edge is deleted from* $succ^{\delta(v)}$.

Lemma 5. *Upon termination of Algorithm 1 at line 11 or line 12, for every* $i \in \{1, ..., n\}$ *and for every* $v \in V$ *it holds that either* $A^i(v) = \perp$ *or* $A^i(v) = A^{\delta(v)}(v)$.

We finish this section with the correctness theorem.

Theorem 3. *Algorithm 1 terminates and upon termination it holds, for all* i, $1 \leq i \leq n$, *that*
- *if* $A^i(v_0) = 1$ *then* $A^G_{min}(v_0) = 1$ *and*
- *if* $A^i(v_0) \in \{?, 0\}$ *then* $A^G_{min}(v_0) = 0$.

Proof. By Lemma 2 we know that Algorithm 1 terminates. For a fixed worker i, by Lemma 3, it certainly holds that if $A^i(v) = 1$ or $A^i(v) = 0$ then $A^G_{min}(v) = A^i(v)$. To show that if $A^i(v) = ?$ then $A^G_{min}(v) = 0$, we first construct a global assignment B such that

$$B(v) = \begin{cases} 0 & \text{if there is } i \in \{1, \ldots, n\} \text{ such that } A^i(v) = ? \text{ or } A^i(v) = 0 \\ 1 & \text{otherwise.} \end{cases} \tag{2}$$

Next we show that B is a fixed-point assignment of G. For a contradiction, let us assume B is not a fixed-point assignment. This can happen in two cases:
- There is a hyper-edge $e = (v, T)$ such that $B(v) = 0$ and $B(u) = 1$ for all $u \in T$. If $A^i(v) = 0$ for some i, it is a direct contradiction with Lemma 3 Condition 2. Otherwise for some i it must hold that $A^i(v) = ?$. By Lemma 5, we get that $A^i(v) = A^{\delta(v)}(v) = ?$. Therefore according to Lemma 3 Condition 3, there exists a configuration u such that $A^{\delta(v)}(u) = ?$ and e is in the dependency set of u. However, $A^{\delta(v)}(u) = ?$ implies that there exists $u \in T$ such that $B(u) = 0$.

– There is a negation edge $e = (v, u)$ such that $B(v) = 0$, and $A^G_{min}(u) = 0$ and e is not deleted. If $A^i(v) = 0$ for some i, it is again a contradiction with Lemma 3 Condition 2. Otherwise for some i it must hold that $A^i(v) = ?$. Then by Lemma 5 we get that $A^i(v) = A^{\delta(v)}(v) = ?$, which is a contradiction with Lemma 4.

Because B is a fixed-point assignment and A^G_{min} is the minimum fixed-point assignment, we get $A^G_{min} \sqsubseteq B$. Therefore if $A^i(v) = ?$ then by the definition of B we have that $B(v) = 0$ and by $A^G_{min}(v) \leq B(v)$ this implies that $A^G_{min}(v) = 0$. □

As a direct consequence of Theorem 3 we get the following corollary.

Corollary 1. *Algorithm 1 terminates and returns* $A^G_{min}(v_0)$.

4 Implementation and Experiments

The single-core local algorithm (local) and its extension with certain zero propagation (czero), together with the distributed versions of czero with non-shared memory and using MPI running on 4 cores (dist-4) and 32 cores (dist-32) have been implemented in an open-source framework written in `C++`. The implementation is available at `http://code.launchpad.net/~tapaal-dist-ctl/verifypn/paper-dist` and contains also all experimental data. The general tool architecture was instantiated to CTL model checking of Petri nets by providing `C++` code for the initial configuration of the EDG and the successor generator (that for a given configuration outputs all outgoing hyper-edges and negation edges). Optionally, one can also implement a custom-made search strategy or choose it from the predefined ones. In our experiments, we use DFS strategy for both the forward and backward propagation (note that even if each worker in the distributed version runs DFS strategy, depending on the actual order of the request arrivals, this may result in pseudo DFS strategies). The framework also includes a console implementation of the game—the integration into the GUI of the tool TAPAAL is currently under development.

To compare the algorithms, we ran experiments on CTL queries for the Petri nets from MCC'16 [27] on machines with four AMD Opteron 6376 processors, each processor having 16 cores. A 15 GB memory limit per core was enforced for all verification runs. We considered all 322 known Petri net models from the competition, each of them coming with 16 different CTL cardinality queries. As many of these models are either trivial to solve or none of the algorithms are able to provide any answer, we first selected an interesting subset of the models where the slowest algorithm used at least 30 seconds on one of the first three queries and at the same time the fastest algorithm solved all three queries within 30 minutes. This left us with 49 models on which we run all 16 CTL queries (in total 784 executions) with the time limit of 1 hour.

Table 1 shows how many queries were answered by the algorithms and documents that our certain zero algorithm solved 90 more queries than the one by Liu and Smolka. Running the distributed algorithm on 4 cores further solved

Algorithm	Answered Queries	Unique Answers
Liu and Smolka Local, 1 core (local)	475	0
Certain Zero Local, 1 core (czero)	565	3
Distributed Certain Zero Local, 4 cores (dist-4)	619	4
Distributed Certain Zero Local, 32 cores (dist-32)	670	52

Table 1: Answered queries within 1 hour (out of 784 executions)

		Query Number															
	Alg.	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
A	local	160	447	–	158	234	250	199	1	228	343	229	241	233	1	223	1
A	czero	157	453	226	154	229	1	1	1	221	100	227	238	232	1	226	1
A	dist-4	82	224	129	86	158	1	1	1	85	1	116	154	133	1	137	1
A	dist-32	21	67	1	20	45	1	1	1	11	1	33	36	46	1	33	1
B	local	465	444	453	16	1	1	401	1	1030	1	877	490	3	458	459	1
B	czero	452	468	464	16	1	1	1	1	522	1	1	477	3	1	2	1
B	dist-4	119	118	125	6	1	1	1	1	180	1	1	144	3	1	1	1
B	dist-32	23	22	23	1	1	1	1	1	1270	1	1	28	1	1	1	1
C	local	343	1	183	85	1	1	4	180	–	1	25	1	165	1	173	172
C	czero	175	1	172	70	1	1	3	1	333	1	23	1	178	1	1	1
C	dist-4	60	1	63	42	3	1	2	1	87	1	12	1	58	1	1	1
C	dist-32	20	2	15	18	2	3	1	1	21	1	11	1	13	1	1	1
D	local	263	446	243	236	219	23	204	356	235	164	1	231	279	1	1	13
D	czero	1	187	6	228	215	21	188	1	220	1	1	229	257	1	1	11
D	dist-4	1	130	6	130	1	12	103	1	122	1	1	124	189	1	1	7
D	dist-32	1	45	2	35	1	3	27	1	38	1	1	41	61	1	1	2
E	local	95	137	140	136	139	135	130	139	139	144	148	1	1	138	132	134
E	czero	96	143	134	134	137	143	129	134	139	146	141	1	1	137	138	1
E	dist-4	33	53	58	53	147	52	50	57	59	65	79	–	1	52	61	1
E	dist-32	30	14	15	14	1225	15	20	16	17	18	19	–	1	16	16	9

Table 2: Verification time in seconds for selected models A: BridgeAndVehicles-PT-V20P20N10, B: Peterson-PT-3, C: ParamProductionCell-PT-4, D: BridgeAndVehicles-PT-V20P10N10, and E: SharedMemory-PT-000010.

54 more queries and the utilisation of 32 cores allowed us to solve additional 51 queries. The number of unique answers—queries that were solved by a given algorithm but not by any of the remaining three algorithms—clearly shows that adding more workers considerably improves the performance of the distributed algorithm. This is despite the fact that we are solving a P-hard problem [15] and such problems are in general believed not to have efficient parallel algorithms.

In Table 2 we zoom in on a few selected models and show the running time (rounded up to the nearest higher second) for all 16 queries of each model. A dash means running out of resources (time or memory). We can observe a significant positive effect of the certain zero propagation on several queries like A.6, B.7, C.8, D.8 and E.16 and in general a satisfactory performance of this technique. The clear trend with multi-core algorithms is that there is usually a considerable speedup when moving from 1 to 4 cores and a generally nice scaling when we employ 32 cores. Here we can often notice reasonable speedups compared to 1

core certain zero algorithm (A.9, B.1, B.2, B.3, B.12, C.9), sometimes even superlinear speedups like in D.5. On the other hand, occasionally using more cores can actually slowdown the computation like in B.9, E.5 or even E.12 where the distributed algorithms did not find the answer at all. These sporadic anomalies can be explained by the pseudo DFS strategy of the distributed algorithm which means that the answer is either discovered immediately like in D.5 or the workers explore significantly more configurations in a portion of the dependency graph where the answer cannot be concluded from. Nevertheless, these unexpected results are rather rare and the general performance of the distributed algorithms, summarized in Table 1, is compelling.

Finally, we also compare the performance of our verification engine with LoLa, the winner in the CTL category at MCC'16 [27]. We run LoLa on all 784 executions (as summarized for our engines in Table 1) with the same 1 hour timeout and 15 GB memory limit. LoLa provided a conclusive answer in 673 cases and given that it is a sequential tool, it won in the comparison with our sequential czero implementation that solved 565 queries. The reason is that about one third of all the 784 queries are actually equivalent to either true or false and hence they can be answered without any state-space exploration by a simple query rewriting technique implemented in LoLa [32]. The problem is that this query simplification implemented in LoLa cannot be turned off for a fair comparison. A detailed analysis revealed that LoLa had 172 exclusive answers compared to 64 exclusive answers of our sequential czero algorithm, however, 58 (34%) of the 172 queries answered exclusively by LoLa were equivalent to either true or false and did not require any state space exploration and further 55 queries (32%) were simplified into a trivial form where LoLa needed to explore less than 1000 markings. After removing these 113 trivial queries, LoLa provided 59 exclusive answers compared to 64 exclusive answers of our sequential czero algorithm. Hence the performance of LoLa is essentially comparable with our sequential algorithm. The main advantage of our approach is that we also provide a distributed implementation that already with 4 cores[3] outperforms the single-core implementation.

5 Conclusion

We extended the formalism of dependency graphs by Liu and Smolka [30] with the notion of negation edges in order to capture nested minimum fixed-point assignments within the same graph. On the extended dependency graphs, we designed an efficient local algorithm that allows us to back-propagate also certain zero values—both along the normal hyper-edges as well as the negation edges and hence considerably speed up the computation. To further increase the performance and applicability of our approach, we suggested to distribute the local algorithm, proved the correctness of the pseudo-code and provided an efficient, open-source implementation. Now the user can take a verification problem, reduce it to an extended dependency graph and get an efficient distributed

[3] The organizers of MCC'17 allow the tools to utilize 4 cores in the competition.

verification engine for free. This is a significant advantage compared to a number of other tools that design a specific distributed algorithm for a fixed modeling language and a fixed property language.

We demonstrated the general applicability of our tool on an example of CTL model checking of Petri nets and evaluated the performance on the benchmark of models from the Model Checking Contest 2016. The results confirm significant improvements over the local algorithm by Liu and Smolka achieved by the certain zero propagation and the distribution of the work among several workers. Already the performance of our sequential algorithm with certain zero propagation is comparable with the world leading tool LoLa for CTL model checking of Petri nets (modulo the query transformation rules implemented additionally in LoLa and not related to the actual state-space search). While LoLa implements only a sequential algorithm, we also provide a generic and efficient distribution of the work among a scalable number of workers.

It was observed that for certain models, the search with a large number of workers can be occasionally directed into a portion of the graph where no conclusive answer can be drawn, implying that sometimes just a few workers find the answer faster. In the future work, we shall look into how to better exploit different search strategies when scaling the number of workers.

Acknowledgments. We would like to thank to Frederik Boenneland, Jakob Dyhr, Mads Johannsen and Torsten Liebke for their help with running LoLa experiments. The work was funded by Sino-Danish Basic Research Center IDEA4CPS, Innovation Fund Denmark center DiCyPS and ERC Advanced Grant LASSO. The last author is partially affiliated with FI MU in Brno.

References

[1] J. Barnat, L. Brim, V. Havel, J. Havlíček, J. Kriho, M. Lenčo, P. Ročkai, V. Štill, and J. Weiser. "DiVinE 3.0 – An Explicit-State Model Checker for Multithreaded C & C++ Programs". In: *Computer Aided Verification (CAV 2013)*. Vol. 8044. LNCS. Springer, 2013, pp. 863–868.

[2] C. Bellettini, M. Camilli, L. Capra, and M. Monga. "Distributed CTL model checking in the cloud". In: *arXiv preprint arXiv:1310.6670* (2013).

[3] B. Bollig, M. Leucker, and M. Weber. "SPIN'02". In: vol. 2318. LNCS. Springer, 2002. Chap. Local Parallel Model Checking for the Alternation-Free μ-Calculus, pp. 128–147.

[4] L. Brim, J. Crhova, and K. Yorav. "Using Assumptions to Distribute CTL Model Checking". In: *ENTCS* 68.4 (2002), pp. 559–574.

[5] F. Cassez, A. David, E. Fleury, K. G. Larsen, and D. Lime. "Efficient on-the-fly algorithms for the analysis of timed games". In: *CONCUR 05*. Vol. 3653. LNCS. Springer, 2005, pp. 66–80.

[6] P. Christoffersen, M. Hansen, A. Mariegaard, J. T. Ringsmose, K. G. Larsen, and R. Mardare. "Parametric Verification of Weighted Systems". In: *SynCoP'15*. Vol. 44. OASIcs. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2015, pp. 77–90.

[7] E. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. "Informatics: 10 Years Back, 10 Years Ahead". In: vol. 2000. LNCS. Springer, 2001. Chap. Progress on the State Explosion Problem in Model Checking, pp. 176–194.

[8] E. M. Clarke and E. A. Emerson. "Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic". In: *Logic of Programs, Workshop*. Springer, 1982, pp. 52–71.

[9] E. M. Clarke, E. A. Emerson, and J. Sifakis. "Model checking: algorithmic verification and debugging". In: *Commun. ACM* 52.11 (2009), pp. 74–84.

[10] A. Dalsgaard, S. Enevoldsen, K. Larsen, and J. Srba. "Distributed Computation of Fixed Points on Dependency Graphs". In: *SETTA'16*. Vol. 9984. LNCS. Springer, 2016, pp. 197–212.

[11] A. David, L. Jacobsen, M. Jacobsen, K. Jørgensen, M. Møller, and J. Srba. "TAPAAL 2.0: Integrated Development Environment for Timed-Arc Petri Nets". In: *TACAS'12*. Vol. 7214. LNCS. Springer, 2012, pp. 492–497.

[12] J. Esparza. "Decidability of model checking for infinite-state concurrent systems". In: *Acta Informatica* 34.2 (1997), pp. 85–107.

[13] H. Garavel, F. Lang, R. Mateescu, and W.Serwe. "CADP 2011: A toolbox for the construction and analysis of distributed processes". In: *STTT* 15.2 (2013), pp. 89–107.

[14] T. Gibson-Robinson, P. Armstrong, A. Boulgakov, and A. Roscoe. "FDR3—A Modern Refinement Checker for CSP". In: *TACAS'14*. Vol. 8413. LNCS. Springer, 2014, pp. 187–201.

[15] R. Greenlaw, H. J. Hoover, and W. L. Ruzzo. *Limits to parallel computation: P-completeness theory*. Vol. 200. Oxford University Press, Inc., 1995.

[16] J. Groote and M. Mousavi. *Modeling and Analysis of Communicating Systems*. The MIT Press, 2014.

[17] O. Grumberg, T. Heyman, and A. Schuster. "Distributed Symbolic Model Checking for μ-Calculus". In: *Formal Methods in System Design* 26.2 (2005), pp. 197–219.

[18] M. Heiner, C. Rohr, and M. Schwarick. "MARCIE–model checking and reachability analysis done efficiently". In: *PN'13*. Vol. 7927. LNCS. Springer, 2013, pp. 389–399.

[19] G. Holzmann. *Spin Model Checker, the: Primer and Reference Manual*. First. Addison-Wesley Professional, 2003.

[20] J. Jensen, K. Larsen, J. Srba, and L. Oestergaard. "Efficient Model Checking of Weighted CTL with Upper-Bound Constraints". In: *STTT* 18.4 (2016), pp. 409–426.

[21] J. F. Jensen, T. Nielsen, L. K. Oestergaard, and J. Srba. "TAPAAL and Reachability Analysis of P/T Nets". In: *ToPNoC XI*. Vol. 9930. Springer, 2016, pp. 307–318.

[22] C. Joubert and R. Mateescu. "Distributed On-the-Fly Model Checking and Test Case Generation". In: *SPIN'06*. Vol. 3925. LNCS. Springer, 2006, pp. 126–145.

[23] G. Kant, A. Laarman, J. Meijer, J. van de Pol, S. Blom, and T. van Dijk. "LTSmin: High-Performance Language-Independent Model Checking". In: *TACAS 2015*. Vol. 9035. LNCS. Springer, 2015, pp. 692–707.

[24] I. Kaufmann, L. S Jensen, and S. Nielsen. "CTL Model Checking of Petri Nets with Dependency Graphs". In: *Semester Project* (2016).

[25] M. Keinänen. *Techniques for Solving Boolean Equation Systems*. Research Report A105. Doctoral dissertation. Helsinki University of Technology, Laboratory for Theoretical Computer Science, 2006, pp. xii+95.

[26] J. J. A. Keiren. "Advanced Reduction Techniques for Model Checking". PhD thesis. Eindhoven University of Technology, 2013.

[27] F. Kordon, H. Garavel, L. M. Hillah, F. Hulin-Hubard, G. Chiardo, A. Hamez, L. Jezequel, A. Miner, J. Meijer, E. Paviot-Adet, D. Racordon, C. Rodriguez, C. Rohr, J. Srba, Y. Thierry-Mieg, G. Trịnh, and K. Wolf. *Complete Results for the 2016 Edition of the Model Checking Contest*. 2016. URL: http://mcc.lip6.fr/2016/results.php.

[28] F. Kordon, H. Garavel, L. M. Hillah, F. Hulin-Hubard, A. Linard, M. Beccuti, A. Hamez, E. Lopez-Bobeda, L. Jezequel, J. Meijer, E. Paviot-Adet, C. Rodriguez, C. Rohr, J. Srba, Y. Thierry-Mieg, and K. Wolf. *Complete Results for the 2015 Edition of the Model Checking Contest*. 2015.

[29] D. Kozen. "ICALP 9". In: vol. 140. LNCS. Springer, 1982. Chap. Results on the propositional μ-calculus, pp. 348–359.

[30] X. Liu and S. A. Smolka. "Simple Linear-Time Algorithms for Minimal Fixed Points". In: *ICALP'98*. Vol. 1443. LNCS. Springer, 1998, pp. 53–66.

[31] L. Tan and R. Cleaveland. "Evidence-Based Model Checking". In: *International Conference on Computer Aided Verification (CAV'02)*. Vol. 2404. LNCS. Springer, 2002, pp. 455–470.

[32] K. Wolf. "Running LoLA 2.0 in a Model Checking Competition". In: *ToPNoC XI*. Vol. 9930. LNCS. Springer, 2016, pp. 274–285.

Appendix

Game Characterizaton

Theorem 1. *Let G be a negation safe EDG, $v \in V$ be a configuration and $r \in \{0,1\}$ be a claim. Then $A^G_{min}(v) = r$ if and only if Defender is the winner of the game starting from the position (v, r).*

Proof. ($\Rightarrow$) Let us first define that a configuration v is of *level* i if v belongs to the component C_i but not to any component C_j where $0 \leq j < i$. By induction on the level of a configuration v, we show that (i) if $A^G_{min}(v) = 0$ then Defender has a winning strategy from $(v, 0)$, and (ii) if $A^G_{min}(v) = 1$ then Defender has a winning strategy from $(v, 1)$.

Let us consider the base case where v is of level 0.

- For the case (i), let us assume that $A^G_{min}(v) = 0$ and consider any play starting from $(v, 0)$. Either Attacker has no outgoing edge v and Defender wins, or for every outgoing hyper-edge (v, T) (notice that there are no negation edges for configurations at level 0) there must be at least one $u \in T$ such that $A^G_{min}(u) = 0$, otherwise A^G_{min} would not be a fixed-point assignment. Defender will choose such u and the play continues from $(u, 0)$. Eventually, either a loop is formed, and the infinite game is winning for Defender as the claim 0 appears infinitely often, or there is no outgoing edge for the attacker to choose, in which case Defender also wins.
- For the case (ii), let us assume that $A^G_{min}(v) = 1$. There must have been a reason why the value of v has been raised from 0 to 1 and the reason is that either v has an outgoing hyper-edge with the empty target set, or there is an outgoing hyper-edge from v such that every node from the target set has the value 1 in the minimum fixed-point assignment. As before, no negation edges can be reached from the component C_0. This means that for the distance function d inductively defined as
  - $d(v) = 0$ if there is a hyper-edge $(v, \emptyset) \in E$, otherwise
  - $d(v) = 1 + \min_{(v,T)\in E} \max_{u \in T} d(u)$,

  we have that $d(v)$ is finite for every v where $A^G_{min}(v) = 1$. Defender's strategy from the position $(v, 1)$ is then to pick from the outgoing hyper-edges (at least one must exist) one that reduces the distance. The distance to the configuration that has a hyper-edge with the empty target set then decreases by at least one (irrelevant of Attacker's choice) and eventually Defender picks such a hyper-edge and Attacker loses the play. Hence Defender has a winning strategy in this case as well.

Let us now consider the inductive case where we have a configuration v of level $i > 0$. Both in the case (i) and (ii) we can now also encounter negation edges.

- For the case (i), Defender still selects configurations from the target set that have the minimum fixed-point value 0, identically with the base case. The only change can be that Attacker can from a configuration v such that $A_{min}(v) = 0$ select also a negation edge $(v, u) \in N$ where $A_{min}(u) = 1$. As the level of u is lower than the level of v, we can use the induction hypothesis to conclude that Defender has a winning strategy from $(u, 1)$.

– For the case (ii), we change the definition of the distance function d such that in the base case $d(v)$ is zero also if there is a negation-edge $(v, u) \in N$ such that $A_{min}(u) = 0$. If the game position becomes such a configuration v, with a negation edge (v, u), then Defender will select that edge and the play continues from $(u, 0)$ that is by induction hypothesis winning for Defender.

Hence the direction from left to right is established.

($\Leftarrow$) We prove the other direction by contraposition. Assume that $A^G_{min}(v) \neq r$ and we want to argue that Defender does not have a universal winning strategy from (v, r) (which by determinacy of the game means that Attacker has a universal winning strategy from (v, r)). However, the fact that $A^G_{min}(v) \neq r$ implies that $A^G_{min}(v) = 1 - r$ and Defender has a winning strategy from $(v, 1 - r)$ as proved above. By the symmetry of the game, this means that Attacker has a winning strategy from (v, r). □

CTL Model Checking of Petri Nets

A *path* in N, starting in a marking M, is a finite or infinite sequence of markings and transition firings, written as

$$M \equiv M_0 \rightarrow M_1 \rightarrow M_2 \rightarrow \dots$$

A path is *maximal* if it is either infinite or ends in a marking M_i such that $M_i \not\rightarrow$; also called a deadlock. The set of all maximal paths for a Petri net N from the marking M is denoted by $\Pi_{max}(M)$.

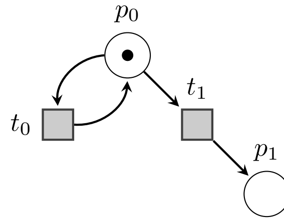

Fig. 4: A Petri net illustrating tokens, places and transitions.

An example of a Petri net is illustrated in Fig. 4. The circles represent places, the rectangles are transitions and arcs that have weight at least one are represented by arrows (in our example all arcs have weight one that we omit this annotation on the arrows). A marking can then be represented as a vector (n_0, n_0) where n_0 denotes the number of tokens in p_0 and n_1 the number of tokens in p_1, respectively. A possible path from the initial marking is $(1, 0)$ is

e.g. $(1,0) \to (1,0) \to (1,0) \to \ldots$. This repeated sequence of markings and firings of the transition t_0 forms an infinite maximal path. Another (finite) maximal path is e.g. $(1,0) \to (1,0) \to (1,0) \to (0,1)$.

Computation Tree Logic

In Computation Tree Logic (CTL), properties are expressed using a combination of logical and temporal operators over a set of basic propositions. In our case the basic propositions express properties of a concrete marking (such as if a certain transition is enabled or the marking contains a certain number of tokens in certain places).

Let $N = (P, T, W, I)$ be a Petri net. To comply with the syntax for the CTL category in MCC'16 [27], we use the following abstract syntax

$$\varphi ::= \mathit{true} \mid \mathit{false} \mid \textbf{is_fireable}(Y) \mid \psi_1 \bowtie \psi_2 \mid \neg\varphi \mid \varphi_1 \wedge \varphi_2 \mid \varphi_1 \vee \varphi_2 \mid$$
$$EG\ \varphi \mid AG\ \varphi \mid EF\ \varphi \mid AF\ \varphi \mid EX\ \varphi \mid AX\ \varphi \mid E\varphi_1 U \varphi_2 \mid A\varphi_1 U \varphi_2$$

$$\psi ::= \psi_1 \oplus \psi_2 \mid c \mid \textbf{token_count}(X)$$

where $\bowtie \in \{<, \leq, =, \geq, >\}$, $X \subseteq P$, $Y \subseteq T$, $c \in \mathbb{N}_0$ and $\oplus \in \{+, -, \cdot\}$. The semantics of a CTL formula φ over a given marking M of the Petri net N is defined in Table 3, using the function $eval_M$ that is given in Table 4. The remaining operators are defined as abbreviations in Table 5.

$M \models \mathit{true}$	
$M \models \neg\varphi$	iff $M \not\models \varphi$
$M \models \varphi_1 \wedge \varphi_2$	iff $M \models \varphi_1$ and $M \models \varphi_2$
$M \models EX\ \varphi$	iff there exists $M' \in M(N)$ where $M \to M'$ and $M' \models \varphi$
$M \models E\varphi_1 U \varphi_2$	iff there exists $(M \equiv M_0 \to M_1 \to M_2 \to \ldots) \in \Pi_{max}(M)$ s.t. there is $i \in \mathbb{N}_0$ where $M_i \models \varphi_2$ and for all $j \in \mathbb{N}_0$ s.t. $0 \leq j < i$ and $M_j \models \varphi_1$
$M \models A\varphi_1 U \varphi_2$	iff for all $(M \equiv M_0 \to M_1 \to M_2 \to \ldots) \in \Pi_{max}(M)$ s.t. there is $i \in \mathbb{N}_0$ where $M_i \models \varphi_2$ and for all $j \in \mathbb{N}_0$ s.t. $0 \leq j < i$ and $M_j \models \varphi_1$
$M \models \textbf{is_fireable}(Y)$	iff there exists $t \in Y$ and M' s.t. $M \xrightarrow{t} M'$
$M \models \psi_1 \bowtie \psi_2$	iff $eval_M(\psi_1) \bowtie eval_M(\psi_2)$

Table 3: CTL Semantics

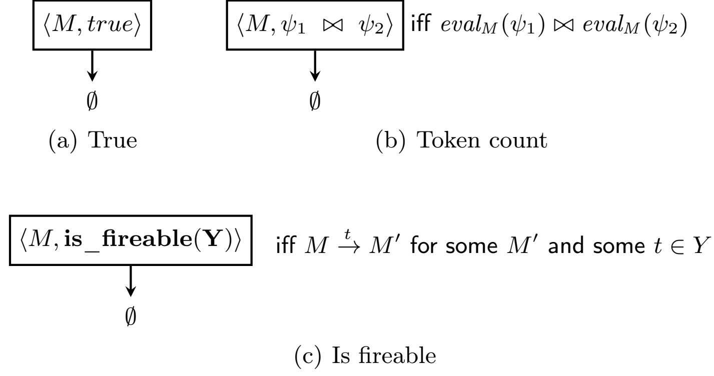

(a) True

(b) Token count

(c) Is fireable

Fig. 5: Atomic rules

$$
\begin{aligned}
&eval_M(c) && = c \\
&eval_M(\textbf{token_count}(X)) && = \textstyle\sum_{p \in X} M(p) \\
&eval_M(e_1 \oplus e_2) && = eval_M(e_1) \oplus eval_M(e_2)
\end{aligned}
$$

Table 4: $eval_M$ semantics

$$
\begin{aligned}
\varphi_1 \vee \varphi_2 &\equiv \neg(\neg\varphi_1 \wedge \neg\varphi_2) \\
AX\ \varphi &\equiv \neg EX\ \neg\varphi \\
EF\ \varphi &\equiv E\ true\ U \varphi \\
AF\ \varphi &\equiv A\ true\ U \varphi \\
EG\ \varphi &\equiv \neg AF\ \neg\varphi \\
AG\ \varphi &\equiv \neg EF\ \neg\varphi \\
false &\equiv \neg true
\end{aligned}
$$

Table 5: Standard abbreviations

Model checking Petri nets using Extended Dependency Graphs

We now reduce the CTL model checking over a Petri net to calculate the minimum fixed point assignment of an EDG. We show how the atomic operators are constructed in Fig. 5. Fig 6 presents the rules for the minimal set of operators required to support the basic formulae from Table 3. Finally in Fig. 7 we also present direct encoding for some of the additional CTL operators. These are

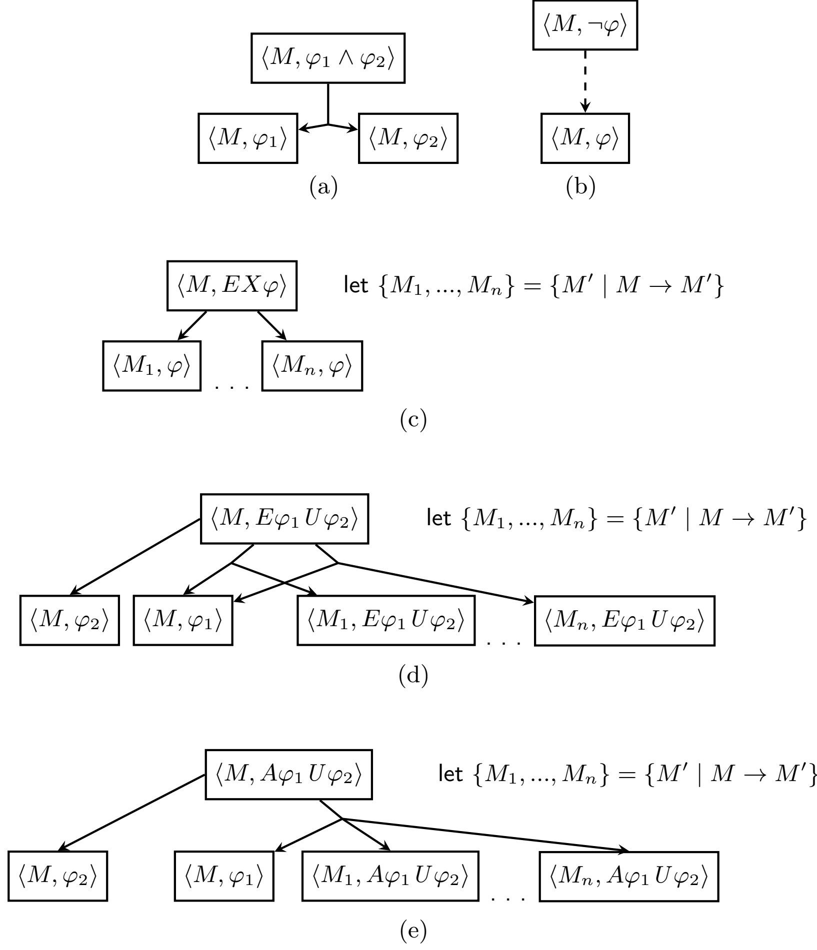

Fig. 6: Minimum set of operators

included in order to limit the amount of configurations required to calculate the minimum fixed point assignment of the extended dependency graph.

In our reduction, each configuration in the extended dependency graph is a pair consisting of a marking M and a CTL formula φ, denoted $\langle M, \varphi \rangle$. Observe that this reduction produces a negation safe EDG.

Theorem 4 (Encoding Correctness). *Let N be a Petri net, M a marking on N and let φ be a CTL formula. Let G be the EDG with the root $\langle M, \varphi \rangle$ constructed as described above. Then $M \models \varphi$ iff $A^G_{min}(\langle M, \varphi \rangle) = 1$.*

The correctness proof of this encoding is available in the report [24].

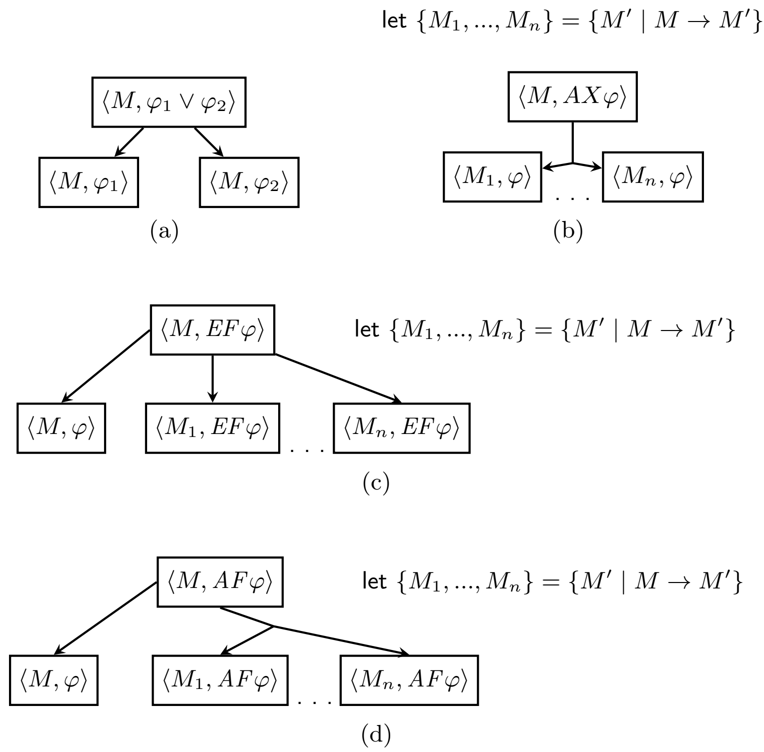

Fig. 7: Extension of operator set

Distributed Algorithm for Fixed-Point Computation

Lemma 1. *During the execution of Algorithm 1, the value of $A^i(v)$ for any worker i and any configuration v will never decrease (with respect to the ordering from Figure 3(c)).*

Proof. First let us observe that the algorithm never assigns $\perp$ to any configuration, hence the only possible way to decrease the assignment value is to assign ? to a configuration which is already assigned 1 or 0. The only place where this can happen is line 2 of the Explore function as the function FinalAssign is always called with only 1 or 0 as an input parameter. However, thanks to the conditions on line 8 of ProcessHyperEdge, line 5 of ProcessNegationEdge and line 7 of ProcessRequest, the Explore function is only called if the previous assignment value is $\perp$. Hence we can never decrease the assignment value of a configuration in any of the local assignments. ☐

Lemma 3. *For any worker i, the repeat-until loop in Algorithm 1 satisfies the following invariants.*

1. *For all $v \in V$, if $A^i(v) = 1$ then $A^G_{min}(v) = 1$.*
2. *For all $v \in V$, if $A^i(v) = 0$ then $A^G_{min}(v) = 0$.*
3. *For all $v \in V$, if $A^i(v) = ?$ and $i = \delta(v)$ then for all $e \in succ^i(v)$ holds that $e \in W^i_E \cup W^i_N$ or $e \in D^i(u)$ for some $u \in V$ where $A^i(u) = ?$.*
4. *For all $v \in V$, if $A^i(v) = ?$ and $i \neq \delta(v)$ then one of the following must hold:*
   - $(v, i) \in M^{\delta(v)}_R$,
   - $i \in C^{\delta(v)}(v)$ *and* $A^{\delta(v)}(v) = ?$, *or*
   - $(v, a) \in M^i_A$ *and* $A^{\delta(v)}(v) = a$ *for some* $a \in \{0, 1\}$.
5. *If there is a negation edge $e = (v, u) \in W^i_N$ s.t. $A^i(u) = ?$ and all workers are idle and v is minimal in all waiting lists and message queues (i.e. for all $(v', x) \in (W_E \cup W_N \cup M_A \cup M_R)$ holds that $dist(v) \leq dist(v')$), then $A^G_{min}(u) = 0$.*

Proof. First we prove Invariants 1 and 2. The only place where the algorithm assigns value 1 or 0 to a configuration is in FinalAssign. Therefore we need to analyse the conditions under which FinalAssign is called. FinalAssign with value 1 or 0 can be called under these circumstances:

- Line 3 of ProcessHyperEdge or line 3 of ProcessNegationEdge where the target is assigned 0. If all targets of a hyper-edge are assigned 1 or the target of a negation edge is assigned 0, it is by the invariant assumption safe to assign 1 also to the source configuration.
- Line 3 of ProcessNegationEdge where the target is assigned ? or 0. The case where the target is 0 is clear thanks to Invariant 2. If the target is assigned ?, this can only happen if the edge was picked based on the fourth condition of PickTask. Therefore the conditions of Invariant 5 apply and it is safe to assign 1 to the source configuration.
- Line 3 of ProcessAnswer. An answer message (a, i) is only sent if $A^{\delta(v)}(v) = a$ and this value is the minimum fixed-point value by Invariants 1 and 2. Therefore it is also safe to assign the same value to $A^i(v)$ in worker i.

– Line 4 of Explore or line 3 of DeleteEdge. If a configuration has no remaining successors that can propagate one, then it is safe to assign 0 to it.

Hence we proved the validity of Invariants 1 and 2.

We shall now focus on Invariant 3. When the value of the assignment is increased from $\perp$ to ? (line 2 of Explore) for a configuration v owned by worker i, all successor edges are pushed into the waiting lists, thus preserving the invariant. By exploring the functions ProcessHyperEdge and ProcessNegationEdge, we observe the following fact. When an edge is picked from the waiting list, one of the following occurs: the source v is assigned a final value, the edge is deleted, or the edge is inserted into the dependency set of some target configuration that is assigned ?. If the target is assigned $\perp$, we call the Explore function that is going to increase it to ?. Finally, when a configuration is assigned 0 or 1, the dependency set is pushed into the waiting lists, therefore the invariant is still preserved.

Let us now discuss Invariant 4. When the value of the assignment is increased from $\perp$ to ? for a configuration v not owned by worker i, the worker sends a request message to the owner (line 7 of Explore), thus the invariant is preserved. As soon as the owner of the configuration receives a request, one of two things happen. If the value of the configuration is already 0 or 1 then the owner sends an answer message to worker i (line 3 of ProcessRequest). Alternatively, if the value of the configuration is $\perp$ or ? then i is inserted into the interested set (line 5 of ProcessRequest) and the value of the configuration is increased from $\perp$ to ? if necessary. Afterwards, when a configuration is assigned 0 or 1, all workers in the interested set are notified via an answer message (line 4 of FinalAssign). Finally, when the answer message is processed by worker i, the configuration is assigned 0 or 1, and the invariant trivially holds too.

We finish by proving Invariant 5. When the conditions of the invariant are satisfied, there are no tasks in any of the waiting and message lists (on any of the workers) that concern the component where the target of the negation edge is located. Since all workers are currently idle, it is also guaranteed that no such task is currently being processed (the opposite would mean that the assignment values in the component can still change as a result of the processing). Therefore it is safe to assume that $A^G_{min}(u) = 0$ as the value of u can never increase to 1, and the invariant holds. □

Lemma 4. *Upon termination of Algorithm 1 at line 11 or line 12, for every negation edge $e = (v, u) \in N$ it holds that either $A^{\delta(v)}(v) \in \{1, \perp\}$ or the negation edge is deleted from $succ^{\delta(v)}$.*

Proof. First, observe that if a negation edge is processed more than once for worker $\delta(v)$, it is either deleted or the source configuration is assigned 1. Hence the target configuration is guaranteed not to be $\perp$. When a negation edge is processed, one of the following will happen:

– the edge is deleted,
– the source configuration is assigned 1, or
– the value of the target configuration is $\perp$. In this case, the edge is re-inserted into the waiting list and will be processed at least twice.

If a negation edge is processed at least once, the condition is satisfied. Observe that if the edge is picked for the first time, and the value of the target configuration is ?, then by Invariant 5, the source configuration can be assigned 1. □

Lemma 5. *Upon termination of Algorithm 1 at line 11 or line 12, for every* $i \in \{1, ..., n\}$ *and for every* $v \in V$ *it holds that either* $A^i(v) = \bot$ *or* $A^i(v) = A^{\delta(v)}(v)$.

Proof. Consider a worker i and a configuration v. If $\delta(v) = i$, the condition holds trivially. If $\delta(v) \neq i$ and $A^i(v) = ?$, then by Lemma 3 Condition 4 also $A^{\delta(v)}(v) = ?$ (since no messages are in transit, because the algorithm has terminated).

If $\delta(v) \neq i$ and $A^i(v) = a \in \{0, 1\}$, it means that worker i at some point received an answer message (v, a). That is because the only place where FinalAssign is called with a configuration that the worker does not own is in ProcessAnswer (and a worker never sends messages to itself). Also, an answer message (v, a) is only sent if the worker who owns v has already assigned it a final value a. Therefore if a worker receives an answer message (v, a) then it is guaranteed that $A^{\delta(v)}(v) = a$. □

Tool Architecture