

Simulation Algorithms for Symbolic Automata (Technical Report)

Lukáš Holík¹, Ondřej Lengál¹, Juraaj Síč^{1,2},
Margus Veanes³, and Tomáš Vojnar¹

¹ FIT, Brno University of Technology, IT4Innovations Centre of Excellence,
Czech Republic

² Faculty of Informatics, Masaryk University, Brno, Czech Republic

³ Microsoft Research, Redmond, USA

Abstract. We investigate means of efficient computation of the simulation relation over symbolic finite automata (SFAs), i.e., finite automata with transitions labeled by predicates over alphabet symbols. In one approach, we build on the algorithm by Ilie, Navaro, and Yu proposed originally for classical finite automata, modifying it using the so-called mintermisation of the transition predicates. This solution, however, generates all Boolean combinations of the predicates, which easily causes an exponential blowup in the number of transitions. Therefore, we propose two more advanced solutions. The first one still applies mintermisation but in a local way, mitigating the size of the exponential blowup. The other one focuses on a novel symbolic way of dealing with transitions, for which we need to sacrifice the counting technique of the original algorithm (counting is used to decrease the dependency of the running time on the number of transitions from quadratic to linear). We perform a thorough experimental evaluation of all the algorithms, together with several further alternatives, showing that all of them have their merits in practice, but with the clear indication that in most of the cases, efficient treatment of symbolic transitions is more beneficial than counting.

1 Introduction

We investigate algorithms for computing simulation relations on states of symbolic finite automata. *Symbolic finite automata* (SFAs) [1,2] extend the classical (nondeterministic) finite automata (NFAs) by allowing one to annotate a transition with a predicate over a possibly infinite alphabet. Such *symbolic* transitions then represent a set of all (possibly infinitely many) concrete transitions over all the individual symbols that satisfy the predicate. SFAs offer a practical solution for automata-based techniques whenever the alphabet is prohibitively large to be processed with a standard NFA, for instance, when processing Unicode-encoded text (e.g., within various security-related analyses) or in automata-based decision procedures for logics such as MSO or WS1S [3,4]. Applications of SFAs over arithmetic alphabets and formulas arise also when dealing with symbolic transducers in the context of various sanitizer and encoder analyses [4].

A *simulation relation* on an automaton underapproximates the inclusion of languages of individual states [20]. This makes it useful for reducing non-deterministic automata and in testing inclusion and equivalence of their languages [20,5,6]. Using simulation for these purposes is often the best compromise between two other alternatives: (i) the cheap but strict bisimulation and (ii) the liberal but expensive language inclusion.

The obvious solution to the problem of computing simulation over an SFA is to use the technique of *mintermisation*: the input SFA is transformed into a form in which predicates on transitions partition the alphabet. Predicates on transitions can then be treated as ordinary alphabet symbols and most of the existing algorithms for NFAs can be used out of the box, including a number of algorithms for computing simulations. We, in particular, consider mintermisation mainly together with the algorithm by Ilie, Navaro, and Yu from [7] (called INY in the following), and, in the experiments, also with the algorithm by Ranzato and Tapparo (called also RT) [19]. A fundamental problem is that mintermisation can increase the number of transitions exponentially due to generating all Boolean combinations of the original transition predicates. Moreover, this problem is not only theoretical, but causes a significant blowup in practice too, as witnessed in the experiments presented in this paper.

We therefore design algorithms that do not need mintermisation. We take as our starting point the algorithm INY, which has the best available time complexity $\mathcal{O}(nm)$ in terms of the number of states n and transitions m of the input NFA. We propose two generalisations of this algorithm. The first one (called LOCALMIN) reflects closely the ideas that INY uses to achieve the low complexity. Instead of applying INY on a globally mintermised SFA, it, however, requires only a *locally mintermised form*: for every state, the predicates on its outgoing transitions partition the alphabet. Local mintermisation is thus exponential only to the maximal out-degree of a state.

Our second algorithm (called NOCOUNT) is fundamentally different from LOCALMIN because it trades off the upfront mintermisation cost against working with predicates in the algorithm, and therefore has a different worst case computational complexity wrt the number of transitions. We show experimentally that this trade-off pays off. To facilitate this trade-off, we had to drop a counting technique that INY uses to improve its time complexity from $\mathcal{O}(n^2m)$ to $\mathcal{O}(nm)$ and that replaces repeated tests for existence of transitions with certain properties by maintaining their number in a dedicated counter and testing it for zero. Dropping the counter-based approach (which depends on at least local mintermisation) in turn allowed an additional optimisation based on aggregating a batch of certain expensive operations (satisfiability checking) on symbolic transitions into one. Overall, this improves the efficiency and ultimately reduces a worst-case 2^m cost, which is typically *independent* of the Boolean algebra, to the cost of inlining the Boolean algebra operations, which may be polynomial or even (sub)linear in m .

In our experiments, although each of the considered algorithms wins in some cases, our new algorithms performed overall significantly better than INY with

global mintermisation. NOCOUNT performed the best overall, which suggests that avoiding mintermisation and aggregating satisfiability tests over transition labels is practically more advantageous than using the counting technique of INY. We have also compared our algorithms with a variant [8] of the RT algorithm, one of the fastest algorithms for computing simulation, run on the globally mintermised automata (we denote the combination as GLOBRT). The main improvement of RT over INY is its use of partition-relation pairs, which allows one to aggregate operations with blocks of the so-far simulation indistinguishable states. Despite this powerful optimisation and the fine-tuned implementation of RT in the VATA library [9], NOCOUNT has a better performance than GLOBRT on automata with high diversity of transition predicates (where mintermisation significantly increases the number of transitions).

Related work. Simulation algorithms for NFAs might be divided between *simple* and *partition-based*. Among the simple algorithms, the algorithm by Henzinger, Henzinger, and Köpke [10] (called HHK) is the first algorithm that achieved the time complexity $\mathcal{O}(nm)$ on Kripke structures. The later algorithm INY [7] is a small modification of HHK and works on finite automata in a time at worst $\mathcal{O}(nm)$. The automata are supposed to be complete (every state has an outgoing transition for every alphabet symbol). INY can be adapted for non-complete automata by adding an initialisation step which costs $\mathcal{O}(\ell n^2)$ time where ℓ is the size of the alphabet, resulting in $\mathcal{O}(nm + \ell n^2)$ overall complexity (cf. §3.1).

The first partition-based algorithm was RT, proposed in [19]. The main innovation of RT is that the overapproximation of the simulation relation is represented by a so-called *partition-relation pair*. In a partition-relation pair, each class of the partition of the set of states represents states that are simulation-equivalent in the current approximation of the simulation, and the relation on the partition denotes the simulation-bigger/smaller classes. Working with states grouped into blocks is faster than working with individual states, and in the case of the most recent partition-based algorithms for Kripke structures [11], it allows to derive the time complexity $\mathcal{O}(n'm)$ where n' is the number of classes of the simulation equivalence (the partition-based algorithms are also significantly faster in practice, although their complexity in terms of m and n is still $\mathcal{O}(nm)$). See e.g. [11] for a more complete overview of algorithms for computing simulation over NFAs and Kripke structures.

Our choice of INY over HHK among the simple algorithms is justified by a smaller dependence of the data structures of INY on the alphabet size. The main reason for basing our algorithms on one of the simple algorithms is their relative simplicity. Partition-based algorithms are intricate as well as the proofs of their small asymptotic complexity. Moreover, they compute predecessors of dynamically refined blocks of states via individual alphabet symbols, which seems to be a problematic step to efficiently generalise for symbolic SFA transitions. Having said that, it remains true that the technique of representing preorders through partition-relation pairs is from the high-level perspective orthogonal to the techniques we have developed to generalise INY. Combining both types of optimisations would be a logical continuation of this work. It is, however, ques-

tionable if generalising already very complex partition-based algorithms, such as [11,19], is the best way to approach computing simulations over SFAs. Most of the intricacy of the partition-based algorithms aims at combining the counting technique with the partition-relation pairs. Our experimental results suggest, however, that rather than using the counting technique, it is more important to optimise the treatment of symbolic transitions and to avoid mintermisation.

Our work complements other works on generalising classical automata algorithms to SFAs, mainly the deterministic minimisation [12] and computing of bisimulation [13].

2 Preliminaries

Throughout the paper, we use the following notation: If $R \subseteq A_1 \times \dots \times A_n$ is an n -ary relation for $n \geq 2$, then $R(x_1, \dots, x_{n-1}) \stackrel{\text{def}}{=} \{y \in A_n \mid R(x_1, \dots, x_{n-1}, y)\}$ for any $x_1 \in A_1, \dots, x_{n-1} \in A_{n-1}$. Let $R^c \stackrel{\text{def}}{=} (A_1 \times \dots \times A_n) \setminus R$.

Effective Boolean algebra. An *effective Boolean algebra* is defined as a tuple $\mathcal{A} = (\mathfrak{D}, \mathbb{P}, \llbracket \cdot \rrbracket, \vee, \wedge, \neg)$ where \mathbb{P} is a set of *predicates* closed under predicate transformers $\vee, \wedge : \mathbb{P} \times \mathbb{P} \rightarrow \mathbb{P}$ and $\neg : \mathbb{P} \rightarrow \mathbb{P}$. A first order interpretation (denotation) $\llbracket \cdot \rrbracket : \mathbb{P} \rightarrow 2^{\mathfrak{D}}$ assigns to every predicate of \mathbb{P} a subset of the *domain* \mathfrak{D} such that, for all $\varphi, \psi \in \mathbb{P}$, it holds that $\llbracket \varphi \vee \psi \rrbracket = \llbracket \varphi \rrbracket \cup \llbracket \psi \rrbracket$, $\llbracket \varphi \wedge \psi \rrbracket = \llbracket \varphi \rrbracket \cap \llbracket \psi \rrbracket$, and $\llbracket \neg \varphi \rrbracket = \mathfrak{D} \setminus \llbracket \varphi \rrbracket$. For $\varphi \in \mathbb{P}$, we write *IsSat*(φ) when $\llbracket \varphi \rrbracket \neq \emptyset$ and say that φ is *satisfiable*. The predicate *IsSat* and the predicate transformers \wedge, \vee , and \neg must be effective (computable). We assume that \mathbb{P} contains predicates \top and \perp with $\llbracket \top \rrbracket = \mathfrak{D}$ and $\llbracket \perp \rrbracket = \emptyset$. Let Φ be a subset of \mathbb{P} . If the denotations of any two distinct predicates in Φ are disjoint, then Φ is called a *partition* (of the set $\bigcup_{\varphi \in \Phi} \llbracket \varphi \rrbracket$). The set *Minterms*(Φ) of *minterms* of a finite set Φ of predicates is defined as the set of all satisfiable predicates of $\{\bigwedge_{\varphi \in \Phi'} \varphi \wedge \bigwedge_{\varphi \in \Phi \setminus \Phi'} \neg \varphi \mid \Phi' \subseteq \Phi\}$. Notice that every predicate of Φ is equivalent to a disjunction of minterms in *Minterms*(Φ).

Below, we assume that it is possible to measure the size of the predicates of the effective Boolean algebra \mathcal{A} that we work with. We denote by $\mathcal{C}_{sat}(x, y)$ the worst-case complexity of constructing a predicate obtained by applying x operations of \mathcal{A} on predicates of the size at most y and checking its satisfiability.

Symbolic finite automata. We define a *symbolic finite automaton* (SFA) as a tuple $M = (Q, \mathcal{A}, \Delta, I, F)$ where Q is a finite set of *states*, $\mathcal{A} = (\mathfrak{D}, \mathbb{P}, \llbracket \cdot \rrbracket, \vee, \wedge, \neg)$ is an effective Boolean algebra, $\Delta \subseteq Q \times \mathbb{P} \times Q$ is a finite *transition relation*, $I \subseteq Q$ is a set of *initial states*, and $F \subseteq Q$ is a set of *final states*. An element (q, ψ, p) of Δ is called a (*symbolic*) *transition* and denoted by $q \{-\psi\} \rightarrow p$. We write $\llbracket q \{-\psi\} \rightarrow p \rrbracket$ to denote the set $\{q \{-a\} \rightarrow p \mid a \in \llbracket \psi \rrbracket\}$ of *concrete transitions* represented by $q \{-\psi\} \rightarrow p$, and we let $\llbracket \Delta \rrbracket \stackrel{\text{def}}{=} \bigcup_{q \{-\psi\} \rightarrow p \in \Delta} \llbracket q \{-\psi\} \rightarrow p \rrbracket$. For $q \in Q$ and $a \in \mathfrak{D}$ let $\Delta(q, a) \stackrel{\text{def}}{=} \{p \in Q \mid q \{-a\} \rightarrow p\}$.

In the following, it is assumed that predicates of all transitions of an SFA are satisfiable unless stated otherwise. A sequence $\rho = q_0 a_1 q_1 a_2 \dots a_n q_n$ with $q_{i-1} \{-a_i\} \rightarrow q_i \in \llbracket \Delta \rrbracket$ for every $1 \leq i \leq n$ is a *run* of M over the word $a_1 \dots a_n$. The

run ρ is *accepting* if $q_0 \in I$ and $q_n \in F$, and a word is *accepted* by M if it has an accepting run. The *language* $\mathcal{L}(M)$ of M is the set of all words accepted by M .

An SFA M is *complete* iff, for all $q \in Q$ and $a \in \mathfrak{D}$, there is $p \in Q$ with $p \xrightarrow{a} q \in \llbracket \Delta \rrbracket$. An SFA can be completed in a straightforward way: from every state q , we add a transition from q labelled with $\neg \bigvee \{ \varphi \mid \exists p \in Q : q \xrightarrow{\varphi} p \in \Delta \}$ to a new non-accepting sink state, if the disjunction is satisfiable.

An SFA M is *globally mintermised* if the set $\mathbb{P}_\Delta \stackrel{\text{def}}{=} \{ \varphi \in \mathbb{P} \mid \exists p, q : p \xrightarrow{\varphi} q \in \Delta \}$ of the predicates appearing on its transitions is a partition. Every SFA can be made globally mintermised by replacing each $p \xrightarrow{\varphi} q \in \Delta$ by the set of transitions $\{ p \xrightarrow{\omega} q \mid \omega \in \text{Minterms}(\mathbb{P}_\Delta) \wedge \text{IsSat}(\omega \wedge \varphi) \}$ (see e.g. [12] for an efficient algorithm), where $\text{IsSat}(\omega \wedge \varphi)$ is an implementation of the test $\llbracket \omega \rrbracket \subseteq \llbracket \varphi \rrbracket$, because if ω is a minterm of \mathbb{P}_Δ and $\varphi \in \mathbb{P}_\Delta$ then $\llbracket \omega \rrbracket \cap \llbracket \varphi \rrbracket \neq \emptyset$ implies that $\llbracket \omega \rrbracket \subseteq \llbracket \varphi \rrbracket$. Since for a set of predicates Φ , the size of $\text{Minterms}(\Phi)$ is at worst $2^{|\Phi|}$, global mintermisation is exponential in the number of transitions.

A classical (*nondeterministic*) *finite automaton* (NFA) $N = (Q, \Sigma, \Delta, I, F)$ over a finite alphabet Σ can be seen as a special case of an SFA where Δ contains solely transitions of the form $q \xrightarrow{a} r$ s.t. $a \in \Sigma$ and $\llbracket a \rrbracket = \{a\}$ for all $a \in \Sigma$. Below, we will sometimes interpret an SFA $M = (Q, \mathcal{A}, \Delta, I, F)$ as its *syntactic NFA* $N = (Q, \mathbb{P}_\Delta, \Delta, I, F)$ in which the predicates are treated as syntactic objects.

Simulation. Let $M = (Q, \mathcal{A}, \Delta, I, F)$ be an SFA. A relation S on Q is a *simulation* on M if whenever $(p, r) \in S$, then the following two conditions hold: (C1) if $p \in F$, then $r \in F$, and (C2) for all $a \in \mathfrak{D}$ and $p' \in Q$ such that $p \xrightarrow{a} p' \in \llbracket \Delta \rrbracket$, there is $r' \in Q$ such that $r \xrightarrow{a} r' \in \llbracket \Delta \rrbracket$ and $(p', r') \in S$. There exists a unique maximal simulation on M , which is reflexive and transitive. We call it the *simulation (preorder)* on M and denote it by \preceq_M (or \preceq when M is clear from the context). Computing \preceq on a given SFA is the subject of this paper. A simulation that is symmetric is called a *bisimulation*, and *the bisimulation equivalence* is the (unique) largest bisimulation, which is always an equivalence relation.

3 Computing Simulation over SFAs

In this section, we present our new algorithms for computing the simulation preorder over SFAs. We start by recalling an algorithm for computing the simulation preorder on an NFA of Ilie, Navarro, and Yu from [7] (called INY), which serves as the basis for our work. Then, we introduce three modifications of INY for SFAs: (i) GLOBINY, (ii) LOCALMIN, and (iii) NOCOUNT. GLOBINY is merely an application of the mintermisation technique: first globally mintermise the SFA and then use INY to compute the NFA simulation preorder over the result. The main contribution of our paper lies in the other two algorithms, which are subtler modifications of INY that avoid global mintermisation by reasoning about the semantics of transition predicates of SFAs.

Before turning to the different algorithms, we start by explaining how \preceq_M can be computed by an abstract fixpoint procedure and provide the intuition behind how such a procedure can be lifted to the symbolic setting.

Abstract procedure for computing \preceq_M . We start by presenting an *abstract fixpoint procedure* for computing the simulation \preceq_M on an SFA $M = (Q, \mathcal{A}, \Delta, I, F)$. We formulate it using the notion of *minimal nonsimulation* $\not\preceq_M$ (which is a dual concept to the maximal simulation \preceq_M introduced before), defined as the least subset $\not\preceq \subseteq Q \times Q$ s.t. for all $s, t \in Q$, it holds that

$$s \not\preceq t \Leftrightarrow (s \in F \wedge t \notin F) \vee \underbrace{\exists i \in Q. \exists a \in \mathfrak{D}. (s \xrightarrow{a} i \wedge \forall j \in Q. (t \xrightarrow{a} j \Rightarrow i \not\preceq j))}_{(1^*)}. \quad (1)$$

Informally, s cannot be simulated by t iff (line 1) s is accepting and t is not, or (line 2) s can continue over some symbol a into i , while t cannot simulate this move by any of its successors j . It is easy to see that $\preceq_M = \not\preceq_M^c$. The algorithms for computing simulation over NFAs are efficient implementations of such a fixpoint procedure using *counter-based* implementations for evaluating (1*). Namely, for every symbol a and a pair of states t and i , it keeps count of those states j that could possibly contradict the universally quantified property. The count dropping to zero means that the property holds universally.

Symbolic abstract procedure for computing \preceq_M . When the domain \mathfrak{D} is very large or infinite, then evaluating (1*) directly is infeasible. If $\text{Minterms}(\mathbb{P}_\Delta)$ is exponentially larger than the set \mathbb{P}_Δ , then evaluating (1*) with a ranging over $\text{Minterms}(\mathbb{P}_\Delta)$ may also be infeasible. Instead, we want to utilize the operations of the algebra \mathcal{A} without explicit reference to elements in \mathfrak{D} and without constructing $\text{Minterms}(\mathbb{P}_\Delta)$. The key insight is that condition (1*) is equivalent to

$$\text{IsSat}(\varphi_{si} \wedge \neg \Gamma(t, \not\preceq^c(i))) \quad (2)$$

where, for $t, s, i \in Q$ and $J \subseteq Q$, we define $\varphi_{si} \stackrel{\text{def}}{=} \bigvee_{(s, \psi, i) \in \Delta} \psi$ and $\Gamma(t, J) \stackrel{\text{def}}{=} \bigvee_{j \in J} \varphi_{tj}$, i.e., $\Gamma(t, \not\preceq^c(i))$ is a disjunction of predicates on all transitions leaving t and entering a state that simulates i . Using (2) to compute (1*) in the abstract procedure thus eliminates the explicit quantification over \mathfrak{D} and avoids computation of $\text{Minterms}(\mathbb{P}_\Delta)$. The equivalence between (1*) and (2) holds because, for all $a \in \mathfrak{D}$ and $R \subseteq Q \times Q$, we have

$$a \in \llbracket \neg \Gamma(t, R^c(i)) \rrbracket \Leftrightarrow \neg \exists j (t \xrightarrow{a} j \wedge (i, j) \in R^c) \Leftrightarrow \forall j (t \xrightarrow{a} j \Rightarrow (i, j) \in R).$$

The fixpoint computation based on (2) is used in our algorithm `NOCOUNT`, which does not require mintermisation. Its disadvantage is that it is not compatible with the counting technique. Our algorithm `LOCALMIN` is then a compromise between mintermisation and `NOCOUNT` that retains the counting technique for the price of using a cheaper, local variant of mintermisation.

3.1 Computing Simulation over NFAs (INY)

In Algorithm 1, we give a slightly modified version of the algorithm `INY` from [7] for computing the simulation preorder over an NFA $N = (Q, \Sigma, \Delta, I, F)$. The

Algorithm 1: INY

Input: An NFA $N = (Q, \Sigma, \Delta, I, F)$
Output: The simulation preorder \preceq_N

- 1 **for** $p, q \in Q, a \in \Sigma$ **do** $N_a(q, p) := |\Delta(q, a)|$;
- 2 $Sim := Q \times Q$;
- 3 $NotSim := F \times (Q \setminus F) \cup \{(q, r) \mid \exists a \in \Sigma : \Delta(q, a) \neq \emptyset \wedge \Delta(r, a) = \emptyset\}$;
- 4 **while** $NotSim \neq \emptyset$ **do**
- 5 remove some (i, j) from $NotSim$ and Sim ;
- 6 **for** $t \xrightarrow{a} j \in \Delta$ **do**
- 7 $N_a(t, i) := N_a(t, i) - 1$;
- 8 **if** $N_a(t, i) = 0$ **then** // $t^{\sigma}i = \emptyset$
- 9 **for** $s \xrightarrow{a} i \in \Delta$ s.t. $(s, t) \in Sim$ **do**
- 10 $NotSim := NotSim \cup \{(s, t)\}$;
- 11 **return** Sim ;

algorithm refines an overapproximation Sim of the simulation preorder until it satisfies the definition of a simulation. The set $NotSim$ is used to store pairs of states (i, j) that were found to contradict the definition of the simulation preorder. $NotSim$ is initialised to contain (a) pairs that contradict condition C1 and (b) pairs that cannot satisfy condition C2 regardless of the rest of the relation, as they relate states with incompatible outgoing symbols. All pairs (i, j) in $NotSim$ are subsequently processed by removing (i, j) from Sim and propagating the change of Sim according to condition C2: for all transitions $t \xrightarrow{a} j \in \Delta$, it is checked whether j was the last a -successor of t that could be simulation-greater than i (hence there are no more such transitions after removing (i, j) from Sim). If this is the case, then t cannot simulate any a -predecessor s of i , and so all such pairs $(s, t) \in Sim$ are added to $NotSim$. In order to have the previous test efficient (a crucial step for the time complexity of the algorithm), the algorithm uses a three-dimensional array of counters $N_a(t, i)$, whose invariant at line 5 is $N_a(t, i) = |t^{\sigma}i|$ where $t^{\sigma}i$ is the set $\Delta(t, a) \cap Sim(i)$ of successors of t over a that simulate i in the current simulation approximation Sim . In order to test $t^{\sigma}i = \emptyset$ —i.e. the second conjunct of (1*)—, it is enough to test if $N_a(t, i) = 0$.

The lemma below shows the time complexity of INY in terms of $n = |Q|$, $m = |\Delta|$, and $\ell = |\Sigma|$. The original paper [7] proves the complexity $O(nm)$ for complete automata, in which case $m \geq \ell n$, so the factor ℓn^2 is subsumed by nm . Since completion of NFAs can be expensive, the initialization step on line 3 of our algorithm is modified (similarly as in [14]) to start with considering states with different sets of symbols appearing on their outgoing transitions as simulation-different; the cost of this step is subsumed by the factor ℓn^2 (see Appendix A for the proof of our formulation of the algorithm).

Lemma 1. INY computes \preceq_N in time $\mathcal{O}(nm + \ell n^2)$.

3.2 Global Mintermisation-based Algorithm for SFAs (GLOBINY)

The algorithm GLOBINY (Algorithm 2) is the initial solution for the problem of computing the simulation preorder over SFAs. It first globally mintermises the

Algorithm 2: GLOBINY

Input: An SFA $M = (Q, \mathcal{A}, \Delta, I, F)$
Output: The simulation preorder \preceq_M
1 $\Delta_G :=$ globally mintermised Δ ;
2 **return** INY($(Q, \mathbb{P}_{\Delta_G}, \Delta_G, I, F)$);

input automaton $M = (Q, \mathcal{A}, \Delta, I, F)$, then interprets the result as an NFA over the alphabet of the minterms, and runs INY on the NFA. The following lemma (together with Lemma 1) implies the correctness of this approach.

Lemma 2. *Let $N = (Q, \mathbb{P}_\Delta, \Delta, I, F)$ be the syntactic NFA of a globally mintermised SFA $M = (Q, \mathcal{A}, \Delta, I, F)$. Then $\preceq_M = \preceq_N$.*

The lemma below shows the time complexity of GLOBINY in terms of $n = |Q|$, $m = |\Delta|$, and the size k of the largest predicate used in Δ .

Lemma 3. *GLOBINY computes \preceq_M in time $\mathcal{O}(nm2^m + \mathcal{C}_{sat}(m, k)2^m)$.*

Intuitively, the complexity follows from the fact each transition of Δ can be replaced by at most 2^m transitions in Δ_G since there can be at most 2^m minterms in $\text{Minterms}(\mathbb{P}_\Delta)$. Nevertheless, 2^m minterms will always be generated (some of them unsatisfiable, though), each of them generated from m predicates of size at most k . More details are available in Appendix B.

3.3 Local Mintermisation-based Algorithm for SFAs (LOCALMIN)

Our next algorithm, called LOCALMIN (Algorithm 3), represents an attempt of running INY on the original SFA without the global mintermisation used above. The main challenge in LOCALMIN is how to symbolically represent the counters $N_a(q, r)$ —representing them explicitly would contradict the idea of symbolic automata and would be impossible if the domain \mathfrak{D} were infinite. We will therefore use counters $N_\psi(q, r)$ indexed with labels ψ of outgoing transitions of q to represent all counters $N_a(q, r)$, with $a \in \llbracket \psi \rrbracket$. A difficulty here is that if the automaton is not globally mintermised, then for some $q \{-\varphi\} \rightarrow p$ and $a, b \in \llbracket \varphi \rrbracket$, the sizes of $q \overset{a}{\rightarrow} r$ and $q \overset{b}{\rightarrow} r$ may differ and hence cannot be represented by a single counter.⁴ For example, if the only outgoing transition of q other than $q \{-\varphi\} \rightarrow p$ is $q \{-\psi\} \rightarrow r$ with $(p, r) \in \text{Sim}$, $\llbracket \varphi \rrbracket = \{a, b\}$, and $\llbracket \psi \rrbracket = \{b\}$, then $|q \overset{a}{\rightarrow} r| = 1$ while $|q \overset{b}{\rightarrow} r| = 2$. To avoid this problem, we introduce the so-called local mintermised form, in which only labels on outgoing transitions of every state must form a partition.

Formally, we say that an SFA $M = (Q, \mathcal{A}, \Delta, I, F)$ is *locally mintermised* if for every state $p \in Q$, the set $\mathbb{P}_{\Delta, p} \stackrel{\text{def}}{=} \{\varphi \in \mathbb{P} \mid \exists q : p \{-\varphi\} \rightarrow q \in \Delta\}$ of the predicates used on the transitions starting from p is a partition. A locally mintermised form is obtained by replacing every transition $p \{-\varphi\} \rightarrow q$ by the set of transitions $\{p \{-\omega\} \rightarrow q \mid \omega \in \text{Minterms}(\mathbb{P}_{\Delta, p}) \wedge \text{IsSat}(\omega \wedge \varphi)\}$. Local mintermisation can hence

⁴ When describing an algorithm that works over an SFA, we use the notation $q \overset{a}{\rightarrow} r$ to represent the set $\llbracket \Delta \rrbracket(q, a) \cap \text{Sim}(r)$, i.e., it refers to the *concrete* transitions of $\llbracket \Delta \rrbracket$.

Algorithm 3: LOCALMIN

Input: A complete SFA $M = (Q, \mathcal{A}, \Delta, I, F)$
Output: The simulation preorder \preceq_M

- 1 $\Delta_L :=$ locally mintermised form of Δ ;
- 2 **for** $p, q \in Q, q \xrightarrow{\psi} t \in \Delta_L$ **do**
- 3 | $N_\psi(q, p) := |\Delta_L(q, \psi)|$;
- 4 $Sim := Q \times Q$; $NotSim := F \times (Q \setminus F)$
- 5 **while** $NotSim \neq \emptyset$ **do**
- 6 | remove some (i, j) from $NotSim$ and Sim ;
- 7 | **for** $t \xrightarrow{\psi_{tj}} j \in \Delta_L$ **do**
- 8 | | $N_{\psi_{tj}}(t, i) := N_{\psi_{tj}}(t, i) - 1$;
- 9 | | **if** $N_{\psi_{tj}}(t, i) = 0$ **then** // $t \xrightarrow{\psi_{tj}} i = \emptyset$
- 10 | | | **for** $s \xrightarrow{\varphi_{si}} i \in \Delta$ s.t. $(s, t) \in Sim$ **do**
- 11 | | | | **if** $IsSat(\psi_{tj} \wedge \varphi_{si})$ **then**
- 12 | | | | | $NotSim := NotSim \cup \{(s, t)\}$;
- 13 **return** Sim ;

be considerably cheaper than global mintermisation as it is only exponential to the maximum out-degree of a state (instead of the number of transitions of the whole SFA). The key property of a locally mintermised SFA M_L is the following: for any transition $q \xrightarrow{\varphi} p$ of M_L and a state $r \in Q$, and for any value of Sim , it holds that $|q^a r|$ is the same for all $a \in \llbracket \varphi \rrbracket$. This means that the set of counters $\{N_a(q, r) \mid a \in \llbracket \varphi \rrbracket\}$ for all symbols in the semantics of φ can be represented by a single counter $N_\varphi(q, r)$.

The use of only locally mintermised transitions also necessitates a modification of the **for** loop on line 6 of INY. In particular, the test on line 9 of INY, which determines the states s that cannot simulate t over the symbol a , only checks syntactic equivalence of the symbols. This could lead to incorrect results because (syntactically) different local minterms of different source states t and s can still have overlapping semantics. It can, in particular, happen that if a counter $N_{\psi_{tj}}(t, i)$, for some predicate ψ_{tj} , reaches zero on line 9 of LOCALMIN, there is a transition from state s to i over a predicate φ_{si} different from ψ_{tj} but with some symbol $a \in \llbracket \varphi_{si} \rrbracket \cap \llbracket \psi_{tj} \rrbracket$. Because of a , the state t cannot simulate s , but this would not happen if the two predicates were only compared syntactically. LOCALMIN solves this issue on lines 10 and 11, where it iterates over all transitions entering i and leaving a state s simulated by t (wrt Sim), and tests whether the predicate φ_{si} on the transition semantically intersects with ψ_{tj} .

LOCALMIN is correct only if the input SFA is complete. As mentioned in §2, this is, however, not an issue, since completion of an SFA is, unlike for NFAs, straightforward, and its cost is negligible compared with the complexity of LOCALMIN presented below.

The lemma below shows the time complexity of LOCALMIN in terms of $n = |Q|$, $m = |\Delta|$, the size k of the largest predicate used in Δ , the *out-degree* m_q for each $q \in Q$ (i.e. the number of transitions leaving q), and the overall *maximum out-degree* $W = \max\{m_q \mid q \in Q\}$.

Algorithm 4: NOCOUNT

Input: A complete SFA $M = (Q, \mathcal{A}, \Delta, I, F)$
Output: The simulation preorder \preceq_M

```
1  $Sim := Q \times Q; NotSim := F \times (Q \setminus F);$   
2 while  $\exists i \in Q : NotSim(i) \neq \emptyset$  do  
3    $Rm := \{t \mid t \rightarrow NotSim(i)\};$   
4    $Sim(i) := Sim(i) \setminus NotSim(i);$   
5    $NotSim(i) := \emptyset;$   
6   for  $t \in Rm$  do  
7      $\psi := \Gamma(t, Sim(i));$   
8     for  $s \xrightarrow{\varphi_{si}} i \in \Delta$  s.t.  $(s, t) \in Sim$  do  
9       if  $IsSat(\neg\psi \wedge \varphi_{si})$  then  
10       $NotSim := NotSim \cup \{(s, t)\};$   
11 return  $Sim;$ 
```

Lemma 4. LOCALMIN derives \preceq_M in time

$$\mathcal{O}\left(n \sum_{q \in Q} m_q 2^{m_q} + m \mathcal{C}_{sat}(W, k) \sum_{q \in Q} 2^{m_q}\right).$$

As shown in more detail in Appendix C, the result can be proved in a similar way as in the case of INY and GLOBINY, taking into account that each transition is, again, replaced by its mintermised versions. This time, however, the mintermised versions are computed independently and locally for each state (and the complexities are summed). Consequently, the factor 2^m gets replaced by 2^{m_q} for the different states $q \in Q$ (together with the replacement of $\mathcal{C}_{sat}(m, k)$ by $\mathcal{C}_{sat}(W, k)$), which can significantly decrease the complexity. On the other hand, as mintermisation is done separately for each state (which can sometimes lead to re-doing some work done only once in GLOBINY) and as one needs the satisfiability test on line 11 of LOCALMIN instead of the purely syntactic test on line 9 of INY, on which GLOBINY is based, GLOBINY can sometimes win in practice. This fact shows up even in our experiments presented in §4.

3.4 Counter-Free Algorithm for SFAs (NOCOUNT)

Before we state our last algorithm, named NOCOUNT (Algorithm 4), let us recall that given an SFA $M = (Q, \mathcal{A}, \Delta, I, F)$, a set $S \subseteq Q$, and a state $q \in Q$, we use $\Gamma(q, S)$ to denote the disjunction of all predicates that reach S from q . We will also write $q \rightarrow S$ to denote that there is a transition from q to some state in S .

In NOCOUNT, we sacrifice the counting technique in order to avoid the local mintermisation (which is still a relatively expensive operation). The obvious price for dropping the counters and local mintermisation is that the emptiness of $t^a i$ for symbols $a \in \psi_{ti}$ can no more be tested in a constant time by asking whether $N_{\psi_{ti}}(t, i) = 0$ as on line 9 of LOCALMIN. It does not even hold any more that $t^a i$ is uniformly empty or non-empty for all $a \in \psi_{ti}$. To resolve the issue, we replace the test from line 9 of LOCALMIN by computing the formula $\psi = \Gamma(t, Sim(i))$

on line 7 of NOCOUNT, which is then used in the test on line 9. Intuitively, ψ represents all b 's such that $t^\psi i$ is *not* empty. By taking the negation of ψ , the test on line 9 of NOCOUNT then explicitly asks whether there is some $a \in \llbracket \varphi_{si} \rrbracket$ for which s can go to i and t cannot simulate this move.

Further, notice that NOCOUNT uses the set Rm for the following optimisation. Namely, if the use of Rm were replaced by an analogy of line 6 from LOCALMIN, it could choose a sequence of several $j \in Q$ such that $(i, j) \in \text{NotSim}$, and then the same ψ would be constructed for each j and tested against the same φ_{si} . In contrast, due to its use of Rm , NOCOUNT will process all $j \in \text{NotSim}(i)$ in a single iteration of the main **while** loop, in which ψ is computed and tested against φ_{si} only once.

Lemma 5 shows the complexity of NOCOUNT in the terms used in Lemma 4.

Lemma 5. NOCOUNT computes \preceq_M in time $\mathcal{O}(n \sum_{q \in Q} m_q^2 + m^2 \mathcal{C}_{sat}(W, k))$.

Observe that $\sum_{q \in Q} m_q = m$ and $W \leq n$, so the above complexity is bounded by $\mathcal{O}(m^2 \mathcal{C}_{sat}(n, k))$. Out-degrees are, however, typically small constants.

The lemma is proved in Appendix D. Compared with the time complexity of LOCALMIN, we can see that, by sacrificing the use of the counters, the complexity becomes quadratic in the number of transitions (since the decrement of the counter on line 8 followed by the test of the counter being zero on line 9 in LOCALMIN is replaced by the computation of Γ on line 7 combined with the test on line 9 in NOCOUNT). On the other hand, since we completely avoid mintermisation, the 2^{m_q} factors are lowered to at most m (m_q in the left-hand side term).

The overall worst-case complexity of NOCOUNT is thus clearly better than those of GLOBINY and LOCALMIN. Moreover, as shown in §4, NOCOUNT is also winning in most of our experiments. Another advantage of avoiding mintermisation is that it often requires a lot of memory. Consequently, GLOBINY and LOCALMIN can run out of memory before even finishing the mintermisation, which is also witnessed in our experiments. If m_q is small for all $q \in Q$ and the predicates do not intersect much, the number of generated minterms can, however, be rather small compared with the number of transitions, and LOCALMIN can in some cases win, as witnessed in our experiments too.

4 Experimental Evaluation

We now present an experimental evaluation of the algorithms from §3 implemented in the Symbolic Automata Toolkit [2]. All experiments were run on an Intel Core i5-3230M CPU@2.6 GHz with 8 GiB of RAM. We used the following two benchmarks:

REGEX. We evaluated the algorithms on SFAs created from 1,921 regular expressions over the UTF-16 alphabet using the \mathbf{BDD}_{16} algebra, which is the algebra of *binary decision diagrams* over 16 Boolean variables representing particular bits of the UTF-16 encoding. These regular expressions were taken from




Fig. 1: Comparison of runtimes of algorithms on SFAs from REGEX. Times are in milliseconds (logarithmic scale).

the website [15], which contains a library of regular expressions created for different purposes, such as matching email addresses, URIs, dates, times, street addresses, phone numbers, etc. The SFAs created from these regular expressions were used before when evaluating algorithms minimising (deterministic) SFAs [12] and when evaluating bisimulation algorithms for SFAs [13]. The largest automaton has 3,190 states and 10,702 transitions; the average transition density of the SFAs is 2.5 transitions per state. Since the UTF-16 alphabet is quite large, a symbolic representation is needed for efficient manipulation of these automata.

WS1S. For this benchmark, we used 131 SFAs generated when deciding formulae of the *weak-monadic second order logic of one successor* (WS1S) [16]. We used two batches of SFAs: 93 deterministic ones from the tool MONA [17] and 38 nondeterministic from DWINA [18]. These automata have at most 2,508 states and 34,374 transitions with the average transition density of 6 transitions per state. These SFAs use the algebra \mathbf{BDD}_k where k is the number of variables in the corresponding formula.

4.1 Comparison of Various Algorithms for Computing Simulation

We first evaluate the effect of our modifications of INY presented in §3. The results presented below clearly show the superiority of our new algorithms over GLOBINY, with NoCOUNT being the overall winner. In addition, we also compare the performance of our new algorithms to a version of the RT algorithm from [19], which is one of the best simulation algorithms. In particular, we use its adaptation for NFAs, which we run after global mintermisation (similarly as INY in GLOBINY). We denote the whole combination GLOBRT. RT is much faster than INY due to its use of the so-called partition-relation pairs to represent the intermediate preorder. Its C++ implementation in the VATA library [9] is also much more optimised than the C# implementation of our algorithms. Despite that, the comparison on automata with many global minterms is clearly favourable to our new algorithms.

To proceed to concrete data, Figs. 1 and 2 show scatter plots of the most interesting comparisons of the runtimes of the considered algorithms on our benchmarks (we give in parentheses the number of times the corresponding algorithm




Fig. 2: Comparison of runtimes of algorithms on SFAs from WS1S. Times are in milliseconds (logarithmic scale).

Table 1: Aggregated results of the performance experiment.

Algorithm	REGEX		WS1S		
	time	wins	time	wins	fails
GLOBINY	12.3 s	2	1,258 s	1	9 (2)
LOCALMIN	11.9 s	0	316 s	0	1 (1)
NoCOUNT	12.4 s	54	44 s	94	0 (0)
GLOBRT	2.8 s	81	594 s	36	3 (2)

won over the other one). The timeout was set to 100 s. Fig. 1 shows the comparison of the algorithms on SFAs from the REGEX benchmark. In this experiment, we removed the SFAs where all algorithms finished within 10 ms (to mitigate the effect of imprecise measurement and noise caused by the C# runtime), which gave us 138 SFAs. Moreover, we also removed one extremely challenging SFA, which dominated the whole benchmark (we report on that SFA, denoted as M_c , later), which left us with the final number of 137 SFAs. On the other hand, Fig. 2 shows the comparison for WS1S. We observe the following phenomena: (i) LOCALMIN is in the majority of cases faster than GLOBINY, (ii) NoCOUNT clearly dominates LOCALMIN, and (iii) the comparison of NoCOUNT and GLOBRT has no clear winner: on the REGEX benchmark, GLOBRT is more often faster, but on the WS1S benchmark, NoCOUNT wins (in many cases, quite significantly).

Further, we also give aggregated results of the experiment in Table 1. In the table, we accumulated the runtimes of the algorithms over the whole benchmark (column “time”) and the number of times each algorithm was the best among all algorithms (column “wins”). The column “fails” shows how many times the respective algorithm failed (by being out of time or memory). In the parentheses, we give the number of times the failure occurred already in the mintermisation. When a benchmark fails, we assign it the time 100 s (the timeout) for the computation of “time”. The times of the challenging SFA M_c from REGEX were: 21 s for GLOBINY, 16 s for LOCALMIN, 25 s for NoCOUNT, and 148 s for GLOBRT. Obviously, including those times would bias the whole evaluation.

Observe that in this comparison, the performance of the algorithms on the two benchmarks differs—although GLOBRT wins on the REGEX benchmark and

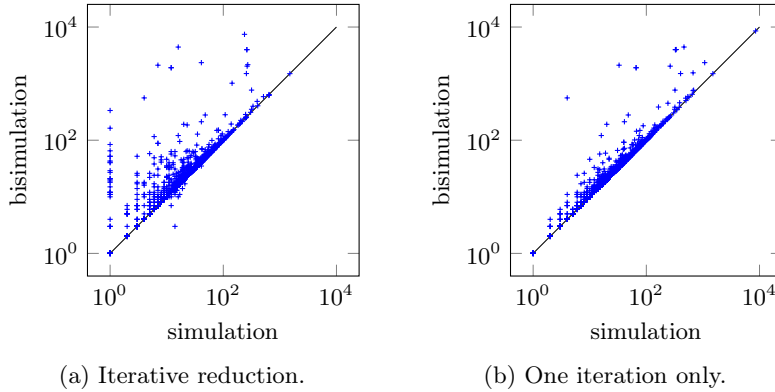


Fig. 3: Simulation vs. bisimulation-based reduction: the number of transitions of the reduced automaton.

the other three algorithms have a comparable overall time (but NOCOUNT still wins in the majority of SFAs among the three), on the more complex benchmark (WS1S), NOCOUNT is the clear winner. The distinct results on the two benchmark sets can be explained by a different diversity of predicates used on the transitions of SFA. In the REGEX benchmark, the globally mintermised automaton has on average 4.5 times more transitions (with the ratio ranging from 1 to 13), while in the WS1S benchmark, the mintermised automaton has on average 23.5 times more transitions (with the ratio ranging from 1 to 716). This clearly shows that our algorithms are effective in avoiding the potential blow-up of mintermisation. As expected, they are slower than RT on examples where the mintermisation is cheap since they do not use the partition-relation data structure.

4.2 Comparison of Simulation and Bisimulation

In the second experiment, we evaluate the benefit of computing simulation over computing bisimulation (we use the implementation of bisimulation computation from [13]). In particular, we focus on an application of (bi-)simulation for (language-preserving) reduction of SFAs from the whole REGEX benchmark.

For every SFA M from the benchmark, we compute its simulation pre-order \preceq_M , take its biggest symmetric fragment (which constitutes an equivalence), and for each of its classes, merge all states of the class into a single state. We also eliminate simulation subsumed transitions (the so-called *little brothers*) using the technique introduced in [20]. In particular, for a state q s.t. there exist transitions $q \xrightarrow{a} p$ and $q \xrightarrow{a} p'$ with $p \preceq p'$, we remove the transition $q \xrightarrow{a} p$ (and also the states that have become unreachable). After that, we reverse the automaton and repeat the whole procedure. These steps continue until the number of states no longer decreases. Similar steps apply to bisimulation (with the exception of taking the symmetric fragment and removing transitions as a bisimulation is already an equivalence).

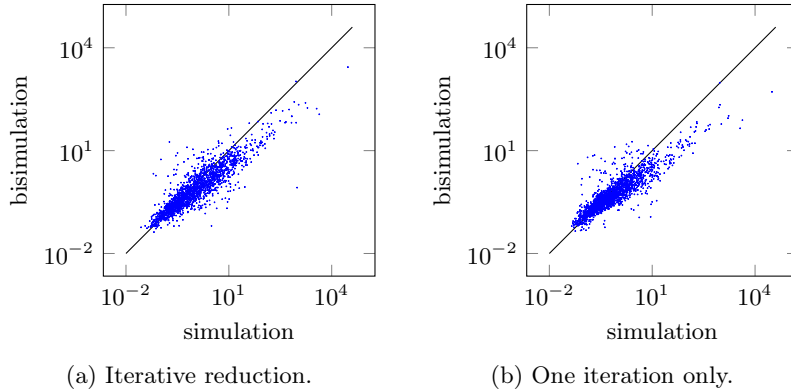


Fig. 4: Simulation vs. bisimulation-based reduction: runtime in milliseconds.

The results comparing the number of transitions of the output SFAs are given in Fig. 3a, showing that the simulation-based reduction is usually much more significant.⁵ Fig. 3b shows the reduction after the first iteration (it corresponds to the “ordinary” simulation and bisimulation-based reduction).

The comparison of the numbers of states gives a very similar picture as the comparison of the numbers of transitions (cf. Appendix E) but simulation wins by a slightly larger margin when comparing the numbers of transitions. This is probably due to the use of the removal of simulation-subsumed transitions, which does not have a meaningful counterpart when working with bisimulations.

As for the runtimes, they differ significantly on the different case studies with some of the cases won by the simulation-based reduction process, some by the bisimulation-based reduction, as can be seen in Fig. 4. Fig. 4a shows comparison of runtimes for the whole iterative process, Fig. 4b shows the comparison for the first iteration only—essentially the time taken by computing the simulation preorder or the bisimulation equivalence. One may see that bisimulation is notably cheaper, especially when the automata are growing larger and both algorithms are taking more time (note the logarithmic scale). Computing simulation was, however, faster in surprisingly many cases.

5 Conclusion and Future Work

We have introduced two new algorithms for computing simulation over symbolic automata that do not depend on global mintermisation: one that needs a local and cheaper variant of mintermisation, and one that does not need mintermisation at all. They perform well especially on automata where mintermisation significantly increases the number of transitions. In the future, we would like to

⁵ There are still some cases when bisimulation achieved a larger reduction than simulation, which may seem unintuitive since the largest bisimulation is always contained in the simulation preorder. This may happen, e.g., when a simulation-based reduction disables an (even greater) reduction on the subsequent reversed SFA.

come up with a partition-based algorithm that could run on an SFA without the need of mintermisation. Such algorithm might, but does not necessarily need to, be based on an NFA partition-based algorithm such as RT. Further, we wish to explore the idea of encoding NFAs over finite alphabets compactly as SFAs over a fast Boolean algebra (such as bit-vector encoding of sets) and compare the performance of our algorithms with known NFA simulation algorithms.

Acknowledgements. The work on this paper was supported by the Czech Science Foundation projects 16-17538S and 16-24707Y, the IT4IXS: IT4Innovations Excellence in Science project (LQ1602), and the FIT BUT internal project FIT-S-17-4014.

References

1. B.W. Watson. Implementing and using finite automata toolkits. Cambridge U. Press (1999).
2. M. Veanes and N. Bjørner. Symbolic automata: The toolkit. In *Proc. of TACAS'12*, LNCS 7214, Springer, 2006.
3. M. Veanes. Applications of symbolic finite automata. In *Proc. of CIAA'13*, LNCS 7982, Springer, 2013.
4. L. D'Antoni and M. Veanes. The power of symbolic automata and transducers. In *Proc. of CAV'17*, LNCS 10426, Springer, 2017.
5. P.A. Abdulla, Y. Chen, L. Holík, R. Mayr, and T. Vojnar, T. When simulation meets antichains. In *Proc. of TACAS'10*, LNCS 6015, Springer, 2010.
6. F. Bonchi and D. Pous. Checking NFA equivalence with bisimulations up to congruence. In *Proc. of POPL'13*, ACM, 2013.
7. L. Ilie, G. Navarro, and S. Yu. On NFA reductions. In *Proc. of Theory is Forever*, LNCS 3113, Springer, 2004.
8. L. Holík and J. Šimáček. Optimizing an LTS-Simulation Algorithm . In *Proc. of MEMICS'09*, Masaryk U., 2009.
9. O. Lengál, J. Šimáček, and T. Vojnar. VATA: A library for efficient manipulation of non-deterministic tree automata. In *Proc. of TACAS'12*, LNCS 7214, Springer, 2012.
10. M.R. Henzinger, T.A. Henzinger, and P.W. Kopke. Computing simulations on finite and infinite graphs. In *Proc. of FOCS'95*, IEEE, 1995.
11. G. Cécé. Foundation for a series of efficient simulation algorithms. In *Proc. of LICS'17*, IEEE, 2017.
12. L. D'Antoni and M. Veanes. Minimization of symbolic automata. In *Proc. of POPL'14*, ACM, 2014.
13. L. D'Antoni and M. Veanes. Forward bisimulations for nondeterministic symbolic finite automata. In *Proc. of TACAS'17*, LNCS 10206, Springer, 2017.
14. M. Eberl. Efficient and verified computation of simulation relations on NFAs. Bachelor's thesis, TU Munich, 2012.
15. Regular expression library, <http://regexlib.com/>.
16. H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. Tree automata techniques and applications. 2007.
17. J. Elgaard, N. Klarlund, and A. Møller. MONA 1.x: New techniques for WS1S and WS2S. In *Proc. of CAV'98*, LNCS 1427, Springer, 1998.

18. T. Fiedor, L. Holík, O. Lengál, and T. Vojnar. Nested antichains for WS1S. In *Proc. of TACAS'15*, LNCS 9035, Springer, 2015.
19. F. Ranzato, and F. Tapparo. A new efficient simulation equivalence algorithm. In *Proc. of LICS'07*, IEEE, 2007.
20. D. Bustan and O. Grumberg. Simulation-based minimization. *ACM Trans. Comput. Logic* 4(2), 2003.

A Complexity of the INY Algorithm

If $n = |Q|$ is the number of states, $m = |\Delta|$ is the number of transitions, and $\ell = |\Sigma|$ is the size of the alphabet of an NFA $N = (Q, \Sigma, \Delta, I, F)$, the time complexity of the INY algorithm is $\mathcal{O}(nm + \ell n^2)$ as stated in Lemma 1. As this fact is not immediately obvious, we give a proof of Lemma 1 below, building on [14].

Proof (Lemma 1). The initialization on lines 1–3 is done in $\mathcal{O}(m + \ell n^2)$ time. Since we save in *NotSim* pairs of states that are to be processed, and we save each pair at most once, line 5 is reached at most n^2 times.

Next, the value of the sum of the initial values of all counters can be characterised as follows:

$$\sum_{\substack{i,t \in Q \\ a \in \Sigma}} N_a(t, i) = \sum_{\substack{i,t \in Q \\ a \in \Sigma}} |\Delta(t, a)|.$$

For a fixed transition $t \in Q$, the sum $\sum_{a \in \Sigma} |\Delta(t, a)|$ is equal to the number of transitions going from the state t , and the sum $\sum_{t \in Q} \sum_{a \in \Sigma} |\Delta(t, a)|$ is then equal to the number of all transitions m . Therefore,

$$\sum_{\substack{i,t \in Q \\ a \in \Sigma}} |\Delta(t, a)| = \sum_{i \in Q} \sum_{\substack{t \in Q \\ a \in \Sigma}} |\Delta(t, a)| = \sum_{i \in Q} m = nm.$$

Also, since counters cannot be negative (because they represent the number of states simulating some state), we can now say that line 7 (decrementing the counters) is reached at most nm times.

Now, the only thing left to show is that lines 9–10 are reached at most nm times. For that, we first note that if we fix $i \in Q, a \in \Sigma$ in $N_a(i, t)$, line 9 is reached at most n times (there are n such counters). On the other hand, if we fix t , the **for** loop on lines 9–10 is iterated at most m times. This stems from a similar fact as the argument for the initial sum of the counters: the **for** loop enumerates all states s , such that $s \xrightarrow{a} i$, and summed over all states i and symbols a , it computes its body m times (for a fixed t). If we combine these two facts, lines 9–10 are reached at most nm times.

Overall, we showed that INY runs in $\mathcal{O}(nm + \ell n^2)$ time. □

B Correctness and Complexity of the GLOBINY Algorithm

In this appendix, we provide a proof of Lemma 2, underlying correctness of Algorithm GLOBINY, and then provide a proof of Lemma 3, stating the complexity of the algorithm.

Proof (Lemma 2). We will prove that $R \subseteq Q \times Q$ is a simulation on M iff R is a simulation on N .

Let R be a simulation on M and assume that it is not simulation on N . Then, there must be some $(q, p) \in R$ contradicting the definition of simulation. Since the sets of final states are the same, (q, p) must contradict Condition C2. This means that there is some $q \xrightarrow{\psi} q' \in \Delta$ for which there is no $p \xrightarrow{\psi} p' \in \Delta$ such that $(q', p') \in R$. However, since for all $a \in \mathfrak{D}_{\mathcal{A}}$ there exists exactly one minterm φ for which $a \in \llbracket \varphi \rrbracket$, this would mean that (q, p) contradicts Condition C2 for M , which is a contradiction.

Since M and N have the same transition relation Δ , the other direction is obvious. \square

Proof (Lemma 3). Apart from n being the number of states of M , m being the number of its transitions, and k being the size of the largest predicate used in the transitions of M , let m' be the number of transition in Δ_G . Recall the definition of \mathcal{C}_{sat} from Section 2. As there can be at most 2^m minterms in $Minterms(\mathbb{P}_{\Delta})$ and every minterm is generated from m predicates, we can compute them in $\mathcal{O}(2^m \mathcal{C}_{sat}(m, k))$ time. Computing Δ_G is then done in $\mathcal{O}(2^m \mathcal{C}_{sat}(m, k) + m2^m)$: every transition is replaced by transitions labelled with minterms. As we then run INY on the syntactic NFA $(Q, \mathbb{P}_{\Delta_G}, \Delta_G, I, F)$, we can conclude that Algorithm GLOBINY has complexity $\mathcal{O}(2^m \mathcal{C}_{sat}(m, k) + m2^m + nm')$. Further, we can bound m' by $m2^{m-1}$ because each transition can be at worst replaced by 2^{m-1} transitions: indeed, note that the predicate of each transition occurs in half of the minterms. The complexity of Algorithm GLOBINY is then

$$\mathcal{O}(2^m \mathcal{C}_{sat}(m, k) + nm2^m).$$

\square

C Complexity of the LOCALMIN Algorithm

In this appendix, we examine the time complexity of Algorithm LOCALMIN and prove Lemma 4.

Proof (Lemma 4). For a given state $q \in Q$, let r_q be the number of minterms in $Minterms(\mathbb{P}_{\Delta, q})$ and m_q the number of transitions with the source state q . Using the same reasoning as for global mintermisation, one can show that Δ_L can be computed in time $\mathcal{O}(\sum_{q \in Q} (r_q \mathcal{C}_{sat}(m_q, k) + m_q r_q))$. Further, let r be the number of all local minterms, i.e., $r = \sum_{q \in Q} r_q$, and let m' be the number of transitions in Δ_L . The initialization on lines 2–4 is done in $\mathcal{O}(nr)$ time:

$|\{r \mid q \xrightarrow{\psi}_{M_L} r\}|$ is computed during mintermisation. Using the same reasoning as in NFA simulation, we can say that, initially, the sum of all counters is nm' , and so line 8 is reached at most nm' times. For a fixed i , line 10 is reached r times because there are r counters $N_{\psi_{t_j}}(t, i)$ that can reach zero only once: for each $t \in Q$ there is one counter $N_{\psi}(t, i)$ for each $\psi \in Minterms(\mathbb{P}_{\Delta, t})$. If we now fix t and ψ_{t_j} , lines 11–12 are reached at most m times. All in all, these lines are reached at most rm times. Since ψ_{t_j} is a minterm created from m_t transitions and $nr \leq nm'$, the time complexity of the algorithm is $\mathcal{O}(\sum_{q \in Q} (r_q \mathcal{C}_{sat}(m_q, k) +$

$m_q r_q) + nm' + mr\mathcal{C}_{sat}(W, k)$). Since, for a given $q \in Q$, r_q is bounded by 2^{m_q} (because r_q is the number of minterms) and since m' is bounded by $\sum_{q \in Q} m_q 2^{m_q}$, the final time complexity of Algorithm LOCALMIN is

$$\mathcal{O}\left(n \sum_{q \in Q} m_q 2^{m_q} + m\mathcal{C}_{sat}(W, k) \sum_{q \in Q} 2^{m_q}\right).$$

□

D Correctness and Complexity of the NOCOUNT Algorithm

We first prove correctness of Algorithm NOCOUNT, i.e., the fact that it indeed computes \preceq_M on an SFA M . Then, we establish the complexity of the algorithm stated in Lemma 5.

D.1 Correctness

We prove that NOCOUNT computes \preceq_M . We first prove that the following invariant is preserved by the algorithm for any simulation relation \preceq over Q .

$$NotSim \subseteq Sim^{\mathbb{G}} \subseteq \preceq^{\mathbb{G}}$$

Initially $NotSim = Sim^{\mathbb{G}} = F \times (Q \setminus F)$, and for all $x \in F$ and $y \in Q \setminus F$ we have $x \not\preceq y$ by definition of simulation.

Consider the main iteration of the **while**-loop and assume that the invariant holds at the start of the **while**-loop. We show that it holds after the **for**-loops.

Fix i such that $NotSim(i) \neq \emptyset$ and $t \in Q$ such that $t \rightarrow NotSim(i)$. Let $\psi = \Gamma(t, Sim(i))$, let $a \in \llbracket \neg\psi \rrbracket$, and let $s \in Q$ be such that $s \{-a\} \rightarrow i$. Suppose, by way of contradiction, that $s \preceq t$. Then there exists j such that $t \{-a\} \rightarrow j$ and $i \preceq j$. But, by the definition of ψ and choice of a , it follows that $j \notin Sim(i)$, i.e., $(i, j) \in Sim^{\mathbb{G}}$, and so $i \not\preceq j$ follows from the invariant, which gives us the contradiction. Hence, $s \not\preceq t$ and Sim is updated by removing (s, t) , i.e., $Sim^{\mathbb{G}}$ as well as $NotSim$ gets the new element (s, t) . Thus, the invariant is preserved, and it follows that $\preceq \subseteq Sim$ and, in particular, that $\preceq_M \subseteq Sim$.

We need to show that upon termination $Sim \subseteq \preceq_M$. We define the *nonsimulation relation of k steps* $\not\preceq_k$ over Q by induction over k as follows. We say that s is *k -nonsimulable by t* when $s \not\preceq_k t$.

$$\begin{aligned} \not\preceq_0 &\stackrel{\text{def}}{=} F \times (Q \setminus F) \\ s \not\preceq_{k+1} t &\stackrel{\text{def}}{=} s \not\preceq_k t \vee \exists a \exists i (s \{-a\} \rightarrow i \wedge \forall j (t \{-a\} \rightarrow j \Rightarrow i \not\preceq_k j)) \end{aligned}$$

Then $\not\preceq_M \stackrel{\text{def}}{=} \not\preceq_\kappa$ where κ is such that $\not\preceq_{\kappa+1} = \not\preceq_\kappa$, and $\preceq_M \stackrel{\text{def}}{=} \not\preceq_M^{\mathbb{G}}$. Thus $s \not\preceq_M t$ means that s is k -nonsimulable by t for some $k \geq 0$, or in other words, t cannot k -step simulate s for any $k \geq 0$. It follows that \preceq_M is the unique *maximal simulation* relation of M . This follows by showing that \preceq_M is indeed

a simulation relation, and (by induction over k) that for any simulation relation \preceq , we have $\not\prec_k \subseteq \preceq^G$. Hence $\preceq \subseteq \preceq_M$.

We show that $\preceq_M^G \subseteq \text{Sim}^G$ by showing that $\not\prec_k \subseteq \text{Sim}^G$ by induction over $k \geq 0$. This holds for $k = 0$ due to the initial value of NotSim and the update of Sim . Assume now that $(s, t) \in \not\prec_{k+1} \setminus \not\prec_k$. Then there exists $a \in \mathfrak{D}$ and $i \in Q$ such that $\llbracket \Delta \rrbracket(t, a) \subseteq \not\prec_k(i)$. Consider the first iteration of the **while**-loop in which, for some such a and i with $s \xrightarrow{\{a\}} i \in \llbracket \Delta \rrbracket$, the last element of $\not\prec_k(i) \cap \llbracket \Delta \rrbracket(t, a)$ is added to $\text{Sim}^G(i)$ on line 4. This must indeed eventually happen because (1) since the automaton is complete, it must hold $\llbracket \Delta \rrbracket(t, a)$ is not empty, and (2) due to the induction hypothesis, all elements of $\not\prec_k(i)$ eventually appear in $\text{NotSim}(i)$ to be later removed from $\text{Sim}(i)$ on line 4. Because the last element of $\not\prec_k(i) \cap \llbracket \Delta \rrbracket(t, a)$ is being removed from $\text{Sim}(i)$, then it is in $\text{NotSim}(i)$, and hence $t \rightarrow \text{NotSim}(i)$. The transition t is therefore added to Rm . When t is processed within the **for**-loop on line 6, the satisfiability check of $\neg\psi \wedge \varphi_{si}$ will eventually succeed because, $\llbracket \neg\Gamma(t, \text{Sim}(i)) \rrbracket = \{a \mid \llbracket \Delta \rrbracket(t, a) \subseteq \text{Sim}^G(i)\}$. And since (by IH) $\not\prec_k \subseteq \text{Sim}^G$ it follows that $\llbracket \Delta \rrbracket(t, a) \subseteq \text{Sim}^G(i)$ and so $a \in \llbracket \neg\psi \rrbracket$. Then (s, t) is added to NotSim and deleted from Sim since all entries that are added to NotSim are later deleted from Sim . Since (s, t) was chosen freely, it follows that $\not\prec_{k+1} \subseteq \text{Sim}^G$ (by termination of NOCOUNT).

D.2 Complexity

Finally, we now proceed to a proof of Lemma 5.

Proof (Lemma 5). The initialization is obviously done in time $\mathcal{O}(n^2)$.

Further, the construction of the set Rm is, for a fixed i , done on the whole in m steps. Indeed, we enumerate all transitions going to some $j \in \text{NotSim}(i)$, and when we sum over all $j \in Q$, we get the number of all transitions in M . Hence, for all $i \in Q$, the computation is done in $\mathcal{O}(nm)$ time.

For fixed states $i, t \in Q$, the state t can occur in the set Rm at most m_t times, and constructing ψ on line 7 consists of iterating through all transitions outgoing from t . Therefore, this line is executed in $\mathcal{O}(m_t^2)$ time. Summed over all $i, t \in Q$, we get $\mathcal{O}(n \sum_{t \in Q} m_t^2)$. Since we assume that the time and space complexity of logical operations are the same, we can assume that the disjunction involved in the computation of Γ has constant complexity and take it into account later, during the check on line 9.⁶ The last fact to show is that lines 9–10 are reached at most m^2 times. Let us again fix states $i, t \in Q$. Then, lines 9–10 can be reached at most $m_t m_i^{-1}$ times where m_i^{-1} is the number of transitions going to i . Summing over all states i, t , we get m^2 .

⁶ To be sure that line 9 is reached at least once, before constructing ψ , we can check whether there are any transitions going into i , and if there are not, we continue with the next iteration of the **for** loop. For readability of the algorithm, we have not included this detail into it.

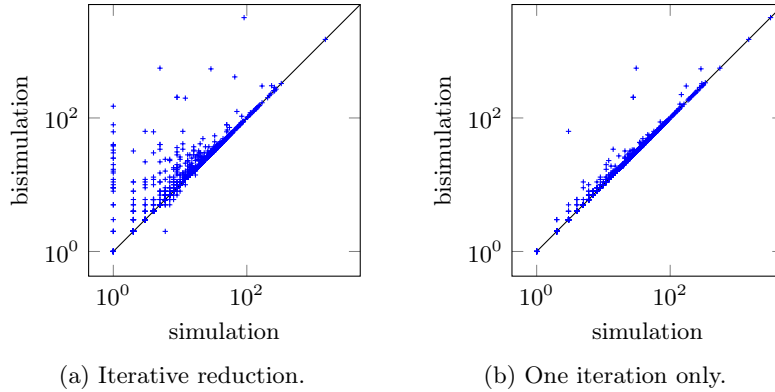


Fig. 5: Simulation vs. bisimulation-based reduction: the number of states of the reduced automaton.

The predicate ψ is a conjunction of m_t predicates, $n^2 \leq nm \leq n \sum_{q \in Q} m_q^2$, and so Algorithm NOCOUNT has the time complexity

$$\mathcal{O}\left(n \sum_{q \in Q} m_q^2 + m^2 \mathcal{C}_{sat}(W, k)\right).$$

□

E Simulation vs. Bisimulation-Based Reduction

We give a more detailed report on the experimental comparison of the effect of simulation and bisimulation based reduction on the REGEX benchmark discussed in Section 4 and also a comparison of the cost of these reductions. Fig. 5 shows a comparison of the numbers of states of the reduced automata. The iterative reducing process described in Section 4 is used on Fig. 5a, Fig. 5b shows the reduction after the first iteration (it corresponds to the “ordinary” simulation and bisimulation-based reduction). Fig. 6 then compares the numbers of transitions. One may see that simulation is clearly more powerful and that it may greatly benefit from iterating the forward and backward reduction. The comparison of the numbers of states gives a very similar picture as the comparison of the numbers of transitions, but one may see that simulation wins by a slightly larger margin when comparing the numbers of transitions. This is probably due to the use of the removal of simulation smaller transitions, which does not have a meaningful counterpart when working with bisimulations.

Lastly, Fig. 7 shows a comparison of the running times of the simulation and bisimulation reduction. Fig. 7a shows the overall time needed by the iterative reduction process, Fig. 7b then the time taken by the first iteration—essentially the time taken by computing the simulation preorder or the bisimulation equivalence. One may see that bisimulation is cheaper overall, especially when the automata are growing larger (note the logarithmic scale). However, computing simulation may be faster in surprisingly many cases.

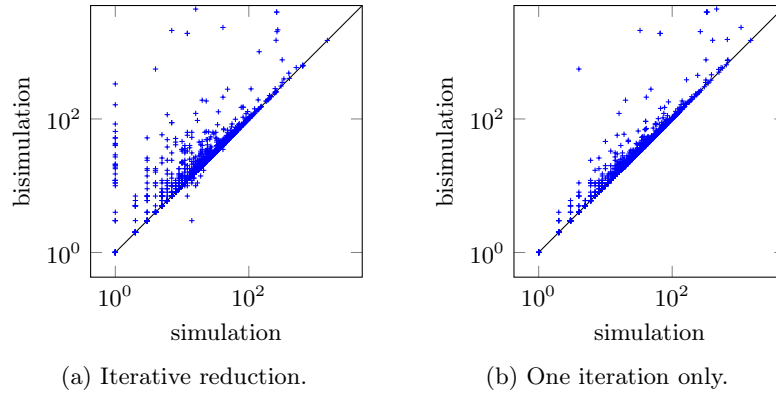


Fig. 6: Simulation vs. bisimulation-based reduction: the number of transitions of the reduced automaton.

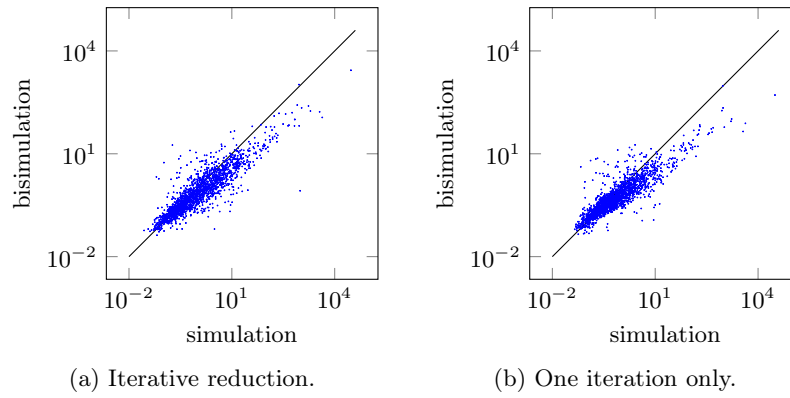


Fig. 7: Simulation vs. bisimulation-based reduction: runtime in milliseconds.