Checking NFA equivalence with bisimulations up to congruence

Filippo Bonchi Damien Pous

CNRS, ENS Lyon, Université de Lyon, LIP (UMR 5668)
{filippo.bonchi,damien.pous}@ens-lyon.fr

Abstract

We introduce *bisimulation up to congruence* as a technique for proving language equivalence of non-deterministic finite automata. Exploiting this technique, we devise an optimisation of the classical algorithm by Hopcroft and Karp [16]. We compare our approach to the recently introduced antichain algorithms, by analysing and relating the two underlying coinductive proof methods. We give concrete examples where we exponentially improve over antichains; experimental results moreover show non negligible improvements.

Keywords Language Equivalence, Automata, Bisimulation, Coinduction, Up-to techniques, Congruence, Antichains.

1. Introduction

Checking language equivalence of finite automata is a classical problem in computer science, which finds applications in many fields ranging from compiler construction to model checking.

Equivalence of deterministic finite automata (DFA) can be checked either via minimisation [9, 15] or through Hopcroft and Karp’s algorithm [2, 16], which exploits an instance of what is nowadays called a *coinduction proof principle* [24, 27, 29]: two states recognise the same language if and only if there exists a *bisimulation* relating them. In order to check the equivalence of two given states, Hopcroft and Karp’s algorithm creates a relation containing them and tries to build a bisimulation by adding pairs of states to this relation: if it succeeds then the two states are equivalent, otherwise they are different.

On the one hand, minimisation algorithms have the advantage of checking the equivalence of all the states at once (while Hopcroft and Karp’s algorithm only check a given pair of states). On the other hand, they have the disadvantage of needing the whole automata from the beginning¹, while Hopcroft and Karp’s algorithm can be executed “on-the-fly” [12], on a lazy DFA whose transitions are computed on demand.

This difference is fundamental for our work and for other recently introduced algorithms based on *antichains* [1, 33]. Indeed, when starting from non-deterministic finite automata (NFA), the

powerset construction used to get deterministic automata induces an exponential factor. In contrast, the algorithm we introduce in this work for checking equivalence of NFA (as well as those in [1, 33]) usually does not build the whole deterministic automaton, but just a small part of it. We write “usually” because in few bad cases, the algorithm still needs exponentially many states of the DFA.

Our algorithm is grounded on a simple observation on deterministic NFA: for all sets X and Y of states of the original NFA, the union (written $+$) of the language recognised by X (written $\llbracket X \rrbracket$) and the language recognised by Y ($\llbracket Y \rrbracket$) is equal to the language recognised by the union of X and Y ($\llbracket X + Y \rrbracket$). In symbols:

$$\llbracket X + Y \rrbracket = \llbracket X \rrbracket + \llbracket Y \rrbracket \quad (1)$$

This fact leads us to introduce a sound and complete proof technique for language equivalence, namely *bisimulation up to context*, that exploits both *induction* (on the operator $+$) and *coinduction*: if a bisimulation R equates both the (sets of) states X_1, Y_1 and X_2, Y_2 , then $\llbracket X_1 \rrbracket = \llbracket Y_1 \rrbracket$ and $\llbracket X_2 \rrbracket = \llbracket Y_2 \rrbracket$ and, by (1), we can immediately conclude that also $X_1 + X_2$ and $Y_1 + Y_2$ are language equivalent. Intuitively, bisimulations up to context are bisimulations which *do not need to relate* $X_1 + X_2$ and $Y_1 + Y_2$ when X_1 (resp. X_2) and Y_1 (resp. Y_2) are already related.

To illustrate this idea, let us check the equivalence of states x and u in the following NFA. (Final states are overlined, labelled edges represent transitions.)

The determinised automaton is depicted below.

Each state is a set of states of the NFA, final states are overlined: they contain at least one final state of the NFA. The numbered lines show a relation which is a bisimulation containing x and u . Actually, this is the relation that is built by Hopcroft and Karp’s algorithm (the numbers express the order in which pairs are added).

The dashed lines (numbered by 1, 2, 3) form a smaller relation which is not a bisimulation, but a bisimulation up to context: the equivalence of states $\{x, y\}$ and $\{u, v, w\}$ could be immediately deduced from the fact that $\{x\}$ is related to $\{u\}$ and $\{y\}$ to $\{v, w\}$, without the need of further exploring the determinised automaton.

Bisimulations up-to, and in particular bisimulations up to context, have been introduced in the setting of concurrency theory [24,

¹ There are few exceptions, like [19] which minimises labelled transition systems w.r.t. bisimilarity rather than trace equivalence.

25, 28] as a proof technique for bisimilarity of CCS or π -calculus processes. As far as we know, they have never been used for proving language equivalence of NFA.

Among these techniques one should also mention *bisimulation up to equivalence*, which, as we show in this paper, is implicitly used in the original Hopcroft and Karp's algorithm. This technique can be briefly explained by noting that not all bisimulations are equivalence relations: it might be the case that a bisimulation relates (for instance) X and Y , Y and Z but not X and Z . However, since $\llbracket X \rrbracket = \llbracket Y \rrbracket$ and $\llbracket Y \rrbracket = \llbracket Z \rrbracket$, we can immediately conclude that X and Z recognise the same language. Analogously to bisimulations up to context, a bisimulation up to equivalence *does not need to relate* X and Z when they are both related to some Y .

The techniques of up-to equivalence and up-to context can be combined resulting in a powerful proof technique which we call *bisimulation up to congruence*. Our algorithm is in fact just an extension of Hopcroft and Karp's algorithm that attempts to build a bisimulation up to congruence instead of a bisimulation up to equivalence. An important consequence, when using up to congruence, is that we do not need to build the whole deterministic automata, but just those states that are needed for the bisimulation up-to. For instance, in the above NFA, the algorithm stops after equating z and $u + v$ and does not build the remaining four states. Despite their use of the up to equivalence technique, this is not the case with Hopcroft and Karp's algorithm, where all accessible subsets of the deterministic automata have to be visited at least once.

The ability of visiting only a small portion of the determinised automaton is also the key feature of the antichain algorithm [33] and its optimisation exploiting similarity [1]. The two algorithms are designed to check *language inclusion* rather than equivalence, but we can relate these approaches by observing that the two problems are equivalent ($\llbracket X \rrbracket = \llbracket Y \rrbracket$ iff $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ and $\llbracket Y \rrbracket \subseteq \llbracket X \rrbracket$; and $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ iff $\llbracket X \rrbracket + \llbracket Y \rrbracket = \llbracket Y \rrbracket$ iff $\llbracket X + Y \rrbracket = \llbracket Y \rrbracket$).

In order to compare with these algorithms, we make explicit the coinductive up-to technique underlying the antichain algorithm [33]. We prove that this technique can be seen as a restriction of up to congruence, for which *symmetry* and *transitivity* are not allowed. As a consequence, the antichain algorithm usually needs to explore more states than our algorithm. Moreover, we show how to integrate the optimisation proposed in [1] in our setting, resulting in an even more efficient algorithm.

Summarising, the contributions of this work are

- (1) the observation that Hopcroft and Karp implicitly use bisimulations up to equivalence (Section 2),
- (2) an efficient algorithm for checking language equivalence (and inclusion), based on a powerful up to technique (Section 3),
- (3) a comparison with antichain algorithms, by recasting them into our coinductive framework (Sections 4 and 5).

Outline

Section 2 recalls Hopcroft and Karp's algorithm for DFA, showing that it implicitly exploits bisimulation up to equivalence. Section 3 describes the novel algorithm, based on bisimulations up to congruence. We compare this algorithm with the antichain one in Section 4, and we show how to exploit similarity in Section 5. Section 6 is devoted to benchmarks. Sections 7 and 8 discuss related and future works. Omitted proofs can be found in the Appendix.

Notation

We denote sets by capital letters $X, Y, S, T, \dots$ and functions by lower case letters $f, g, \dots$. Given sets X and Y , $X \times Y$ is their Cartesian product, $X \uplus Y$ is the disjoint union and X^Y is the set of functions $f: Y \rightarrow X$. Finite iterations of a function $f: X \rightarrow X$

are denoted by f^n (formally, $f^0(x) = x$, $f^{n+1}(x) = f(f^n(x))$). The collection of subsets of X is denoted by $\mathcal{P}(X)$. The (omega) iteration of a function $f: \mathcal{P}(X) \rightarrow \mathcal{P}(X)$ is denoted by f^ω (formally, $f^\omega(Y) = \bigcup_{n \geq 0} f^n(Y)$). For a set of letters A , A^* denotes the set of all finite words over A ; ϵ the empty word; and $w_1 w_2$ the concatenation of words $w_1, w_2 \in A^*$. We use 2 for the set $\{0, 1\}$ and 2^{A^*} for the set of all languages over A .

2. Hopcroft and Karp's algorithm for DFA

A deterministic finite automaton (DFA) over the alphabet A is a triple (S, o, t) , where S is a finite set of states, $o: S \rightarrow 2$ is the output function, which determines if a state $x \in S$ is final ($o(x) = 1$) or not ($o(x) = 0$), and $t: S \rightarrow S^A$ is the transition function which returns, for each state x and for each letter $a \in A$, the next state $t_a(x)$. For $a \in A$, we write $x \xrightarrow{a} x'$ to mean that $t_a(x) = x'$. For $w \in A^*$, we write $x \xrightarrow{w} x'$ for the least relation such that (1) $x \xrightarrow{\epsilon} x$ and (2) $x \xrightarrow{aw'} x'$ iff $x \xrightarrow{a} x''$ and $x'' \xrightarrow{w'} x'$.

For any DFA, there exists a function $\llbracket - \rrbracket: S \rightarrow 2^{A^*}$ mapping states to languages, defined for all $x \in S$ as follows:

$$\llbracket x \rrbracket(\epsilon) = o(x) \quad , \quad \llbracket x \rrbracket(aw) = \llbracket t_a(x) \rrbracket(w) \quad .$$

The language $\llbracket x \rrbracket$ is called the language accepted by x . Given two automata (S_1, o_1, t_1) and (S_2, o_2, t_2) , the states $x_1 \in S_1$ and $x_2 \in S_2$ are said to be *language equivalent* (written $x_1 \sim x_2$) iff they accept the same language.

Remark 1. In the following, we will always consider the problem of checking the equivalence of states of one single and fixed automaton (S, o, t) . We do not lose generality since for any two automata (S_1, o_1, t_1) and (S_2, o_2, t_2) it is always possible to build an automaton $(S_1 \uplus S_2, o_1 \uplus o_2, t_1 \uplus t_2)$ such that the language accepted by every state $x \in S_1 \uplus S_2$ is the same as the language accepted by x in the original automaton (S_i, o_i, t_i) . For this reason, we also work with automata without explicit initial states: we focus on the equivalence of two arbitrary states of a fixed DFA.

2.1 Proving language equivalence via coinduction

We first define bisimulation. We make explicit the underlying notion of progression which we need in the sequel.

Definition 1 (Progression, Bisimulation). *Given two relations $R, R' \subseteq S \times S$ on states, R progresses to R' , denoted $R \succ R'$, if whenever $x R y$ then*

1. $o(x) = o(y)$ and
2. for all $a \in A$, $t_a(x) R' t_a(y)$.

A bisimulation is a relation R such that $R \succ R$.

As expected, bisimulation is a sound and complete proof technique for checking language equivalence of DFA:

Proposition 1 (Coinduction). *Two states are language equivalent iff there exists a bisimulation that relates them.*

2.2 Naive algorithm

Figure 1 shows a naive version of Hopcroft and Karp's algorithm for checking language equivalence of the states x and y of a deterministic finite automaton (S, o, t) . Starting from x and y , the algorithm builds a relation R that, in case of success, is a bisimulation. In order to do that, it employs the set (of pairs of states) *todo* which, intuitively, at any step of the execution, contains the pairs (x', y') that must be checked: if (x', y') already belongs to R , then it has already been checked and nothing else should be done. Otherwise, the algorithm checks if x' and y' have the same outputs (i.e., if both are final or not). If $o(x') \neq o(y')$, then x and y are different.

Naive(x, y)

```

(1)  $R$  is empty;  $todo$  is empty;
(2) insert  $(x, y)$  in  $todo$ ;
(3) while  $todo$  is not empty, do {
  (3.1) extract  $(x', y')$  from  $todo$ ;
  (3.2) if  $(x', y') \in R$  then skip;
  (3.3) if  $o(x') \neq o(y')$  then return false;
  (3.4) for all  $a \in A$ ,
    insert  $(t_a(x'), t_a(y'))$  in  $todo$ ;
  (3.5) insert  $(x', y')$  in  $R$ ;
(4) return true;

```

Figure 1. Naive algorithm for checking the equivalence of states x and y of a DFA (S, o, t) ; R and $todo$ are sets of pairs of states. The code of $HK(x, y)$ is obtained by replacing step 3.2 with `if  $(x', y') \in e(R)$  then skip`.

Figure 2. Checking for DFA equivalence.

If $o(x') = o(y')$, then the algorithm inserts (x', y') in R and, for all $a \in A$, the pairs $(t_a(x'), t_a(y'))$ in $todo$.

Proposition 2. For all $x, y \in S$, $x \sim y$ iff $\text{Naive}(x, y)$.

Proof. We first observe that if $\text{Naive}(x, y)$ returns true then the relation R that is built before arriving to step 4 is a bisimulation. Indeed, the following proposition is an invariant for the loop corresponding to step 3:

$$R \rightsquigarrow R \cup todo$$

This invariant is preserved since at any iteration of the algorithm, a pair (x', y') is removed from $todo$ and inserted in R after checking that $o(x') = o(y')$ and adding $(t_a(x'), t_a(y'))$ for all $a \in A$ in $todo$. Since $todo$ is empty at the end of the loop, we eventually have $R \rightsquigarrow R$, i.e., R is a bisimulation. By Proposition 1, $x \sim y$.

We now prove that if $\text{Naive}(x, y)$ returns false, then $x \not\sim y$. Note that for all (x', y') inserted in $todo$, there exists a word $w \in A^*$ such that $x \xrightarrow{w} x'$ and $y \xrightarrow{w} y'$. Since $o(x') \neq o(y')$, then $\llbracket x' \rrbracket(\epsilon) \neq \llbracket y' \rrbracket(\epsilon)$ and thus $\llbracket x \rrbracket(w) = \llbracket x' \rrbracket(\epsilon) \neq \llbracket y' \rrbracket(\epsilon) = \llbracket y \rrbracket(w)$, that is $x \not\sim y$. $\square$

Since both Hopcroft and Karp's algorithm and the one we introduce in Section 3 are simple variations of this naive one, it is important to illustrate its execution with an example. Consider the DFA with input alphabet $A = \{a\}$ in the left-hand side of Figure 2, and suppose we want to check that x and u are language equivalent.

During the initialisation, (x, u) is inserted in $todo$. At the first iteration, since $o(x) = 0 = o(u)$, (x, u) is inserted in R and (y, v) in $todo$. At the second iteration, since $o(y) = 1 = o(v)$, (y, v) is inserted in R and (z, w) in $todo$. At the third iteration, since $o(z) = 0 = o(w)$, (z, w) is inserted in R and (y, v) in $todo$. At the fourth iteration, since (y, v) is already in R , the algorithm does nothing. Since there are no more pairs to check in $todo$, the relation R is a bisimulation and the algorithm terminates returning true.

These iterations are concisely described by the numbered dashed lines in Figure 2. The line i means that the connected pair is inserted in R at iteration i . (In the sequel, when enumerating iterations, we ignore those where a pair from $todo$ is already in R so that there is nothing to do.)

Remark 2. Unless it finds a counter-example, *Naive* constructs the smallest bisimulation that relates the two starting states (see Proposition 8 in Appendix A). On the contrary, minimisation algorithms [9, 15] are designed to compute the largest bisimulation relation for a given automaton. For instance, taking automaton on the left of Figure 2, they would equate the states x and w which are language equivalent, while $\text{Naive}(x, u)$ does not relate them.

2.3 Hopcroft and Karp's algorithm

The naive algorithm is quadratic: a new pair is added to R at each non-trivial iteration, and there are only n^2 such pairs, where $n = |S|$ is the number of states of the DFA. To make this algorithm (almost) linear, Hopcroft and Karp actually record a set of *equivalence classes* rather than a set of visited pairs. As a consequence, their algorithm may stop earlier, when an encountered pair of states is not already in R but in its reflexive, symmetric, and transitive closure. For instance in the right-hand side example from Figure 2, we can stop when we encounter the dotted pair (y, w) , since these two states already belong to the same equivalence class according to the four previous pairs.

With this optimisation, the produced relation R contains at most n pairs (two equivalence classes are merged each time a pair is added). Formally, and ignoring the concrete data structure to store equivalence classes, Hopcroft and Karp's algorithm consists in simply replacing step 3.2 in Figure 1 with

```

(3.2) if  $(x', y') \in e(R)$  then skip;

```

where $e: \mathcal{P}(S \times S) \rightarrow \mathcal{P}(S \times S)$ is the function mapping each relation $R \subseteq S \times S$ into its symmetric, reflexive, and transitive closure. We hereafter refer to this algorithm as HK.

2.4 Bisimulations up-to

We now show that the optimisation used by Hopcroft and Karp corresponds to exploiting an “up-to technique”.

Definition 2 (Bisimulation up-to). Let $f: \mathcal{P}(S \times S) \rightarrow \mathcal{P}(S \times S)$ be a function on relations on S . A relation R is a bisimulation up to f if $R \rightsquigarrow f(R)$, i.e., whenever $x R y$ then

1. $o(x) = o(y)$ and
2. for all $a \in A$, $t_a(x) f(R) t_a(y)$.

With this definition, Hopcroft and Karp's algorithm just consists in trying to build a bisimulation up to e . To prove the correctness of the algorithm it suffices to show that any bisimulation up to e is contained in a bisimulation. We use for that the notion of compatible function [26, 28]:

Definition 3 (Compatible function). A function $f: \mathcal{P}(S \times S) \rightarrow \mathcal{P}(S \times S)$ is compatible if it is monotone and it preserves progressions: for all $R, R' \subseteq S \times S$,

$$R \rightsquigarrow R' \text{ entails } f(R) \rightsquigarrow f(R').$$

Proposition 3. Let f be a compatible function. Any bisimulation up to f is contained in a bisimulation.

Proof. Suppose that R is a bisimulation up to f , i.e., that $R \rightsquigarrow f(R)$. Using compatibility of f and by a simple induction on n , we get $\forall n, f^n(R) \rightsquigarrow f^{n+1}(R)$. Therefore, we have

$$\bigcup_n f^n(R) \rightsquigarrow \bigcup_n f^n(R),$$

in other words, $f^\omega(R) = \bigcup_n f^n(R)$ is a bisimulation. This latter relation trivially contains R , by taking $n = 0$. $\square$

We could prove directly that e is a compatible function; we however take a detour to ease our correctness proof for the algorithm we propose in Section 3.

Lemma 1. *The following functions are compatible:*

- id : the identity function;
- $f \circ g$: the composition of compatible functions f and g ;
- $\bigcup F$: the pointwise union of an arbitrary family F of compatible functions: $\bigcup F(R) = \bigcup_{f \in F} f(R)$;
- f^ω : the (omega) iteration of a compatible function f .

Lemma 2. *The following functions are compatible:*

- the constant reflexive function: $r(R) = \{(x, x) \mid \forall x \in S\}$;
- the converse function: $s(R) = \{(y, x) \mid x R y\}$;
- the squaring function: $t(R) = \{(x, z) \mid \exists y, x R y R z\}$.

Intuitively, given a relation R , $(s \cup id)(R)$ is the symmetric closure of R , $(r \cup s \cup id)(R)$ is its reflexive and symmetric closure, and $(r \cup s \cup t \cup id)^\omega(R)$ is its symmetric, reflexive and transitive closure: $e = (r \cup s \cup t \cup id)^\omega$. Another way to understand this decomposition of e is to recall that for a given R , $e(R)$ can be defined inductively by the following rules:

$$\frac{}{x e(R) x} r \quad \frac{x e(R) y}{y e(R) x} s \quad \frac{x e(R) y \quad y e(R) z}{x e(R) z} t \quad \frac{x R y}{x e(R) y} id$$

Theorem 1. *Any bisimulation up to e is contained in a bisimulation.*

Proof. By Proposition 3, it suffices to show that e is compatible, which follows from Lemma 1 and Lemma 2. $\square$

Corollary 1. *For all $x, y \in S$, $x \sim y$ iff $HK(x, y)$.*

Proof. Same proof as for Proposition 2, by using the invariant $R \mapsto e(R) \cup todo$. We deduce that R is a bisimulation up to e after the loop. We conclude with Theorem 1 and Proposition 1. $\square$

Returning to the right-hand side example from Figure 2, Hopcroft and Karp's algorithm constructs the relation

$$R_{HK} = \{(x, u), (y, v), (z, w), (z, v)\}$$

which is not a bisimulation, but a bisimulation up to e : it contains the pair (x, u) , whose b -transitions lead to (y, w) , which is not in R_{HK} but in its equivalence closure, $e(R_{HK})$.

3. Optimised algorithm for NFA

We now move from DFA to non-deterministic automata (NFA). We start with standard definitions about semi-lattices, determinisation, and language equivalence for NFA.

A *semi-lattice* $(X, +, 0)$ consists of a set X and a binary operation $+$: $X \times X \rightarrow X$ which is associative, commutative, idempotent (ACI), and has $0 \in X$ as identity. Given two semi-lattices $(X_1, +_1, 0_1)$ and $(X_2, +_2, 0_2)$, an *homomorphism* of semi-lattices is a function f : $X_1 \rightarrow X_2$ such that for all $x, y \in X_1$, $f(x +_1 y) = f(x) +_2 f(y)$ and $f(0_1) = 0_2$. The set $2 = \{0, 1\}$ is a semi-lattice when taking $+$ to be the ordinary Boolean or. Also the set of all languages 2^{A^*} carries a semi-lattice where $+$ is the union of languages and 0 is the empty language. More generally, for any set X , $\mathcal{P}(X)$ is a semi-lattice where $+$ is the union of sets and 0 is the empty set. In the sequel, we indiscriminately use 0 to denote the element $0 \in 2$, the empty language in 2^{A^*} , and the

empty set in $\mathcal{P}(X)$. Similarly, we use $+$ to denote the Boolean or in 2 , the union of languages in 2^{A^*} , and the union of sets in $\mathcal{P}(X)$.

A non-deterministic finite automaton (NFA) over the input alphabet A is a triple (S, o, t) , where S is a finite set of states, o : $S \rightarrow 2$ is the output function (as for DFA), and t : $S \rightarrow \mathcal{P}(S)^A$ is the transition relation, which assigns to each state $x \in S$ and input letter $a \in A$ a set of possible successor states.

The *powerset construction* transforms any NFA (S, o, t) in the DFA $(\mathcal{P}(S), o^\#, t^\#)$ where $o^\#$: $\mathcal{P}(S) \rightarrow 2$ and $t^\#$: $\mathcal{P}(S) \rightarrow \mathcal{P}(S)^A$ are defined for all $X \in \mathcal{P}(S)$ and $a \in A$ as follows:

$$o^\#(X) = \begin{cases} o(x) & \text{if } X = \{x\} \text{ with } x \in S \\ 0 & \text{if } X = 0 \\ o^\#(X_1) + o^\#(X_2) & \text{if } X = X_1 + X_2 \end{cases}$$

$$t_a^\#(X) = \begin{cases} t_a(x) & \text{if } X = \{x\} \text{ with } x \in S \\ 0 & \text{if } X = 0 \\ t_a^\#(X_1) + t_a^\#(X_2) & \text{if } X = X_1 + X_2 \end{cases}$$

Observe that in $(\mathcal{P}(S), o^\#, t^\#)$, the states form a semi-lattice $(\mathcal{P}(S), +, 0)$, and $o^\#$ and $t^\#$ are, by definition, semi-lattices homomorphisms. These properties are fundamental for the up-to technique we are going to introduce; in order to highlight the difference with generic DFA (which usually do not carry this structure), we introduce the following definition.

Definition 4. *A determinised NFA is a DFA $(\mathcal{P}(S), o^\#, t^\#)$ obtained via the powerset construction of some NFA (S, o, t) .*

Hereafter, we use a new notation for representing states of determinised NFA: in place of the singleton $\{x\}$ we just write x and, in place of $\{x_1, \dots, x_n\}$, we write $x_1 + \dots + x_n$.

For an example, consider the NFA (S, o, t) depicted below (left) and part of the determinised NFA $(\mathcal{P}(S), o^\#, t^\#)$ (right).

In the determinised NFA, x makes one single a -transition going into $y + z$. This state is final: $o^\#(y + z) = o^\#(y) + o^\#(z) = o(y) + o(z) = 1 + 0 = 1$; it makes an a -transition into $t_a^\#(y + z) = t_a^\#(y) + t_a^\#(z) = t_a(y) + t_a(z) = x + y$.

The language accepted by the states of a NFA (S, o, t) can be conveniently defined via the powerset construction: the language accepted by $x \in S$ is the language accepted by the singleton $\{x\}$ in the DFA $(\mathcal{P}(S), o^\#, t^\#)$, in symbols $\llbracket \{x\} \rrbracket$. Therefore, in the following, instead of considering the problem of language equivalence of states of the NFA, we focus on language equivalence of *sets* of states of the NFA: given two sets of states X and Y in $\mathcal{P}(S)$, we say that X and Y are language equivalent ($X \sim Y$) iff $\llbracket X \rrbracket = \llbracket Y \rrbracket$. This is exactly what happens in standard automata theory, where NFA are equipped with sets of initial states.

3.1 Extending coinduction to NFA

In order to check if two sets of states X and Y of an NFA (S, o, t) are language equivalent, we can simply employ the bisimulation proof method on $(\mathcal{P}(S), o^\#, t^\#)$. More explicitly, a bisimulation for a NFA (S, o, t) is a relation $R \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ on sets of states, such that whenever $X R Y$ then (1) $o^\#(X) = o^\#(Y)$, and (2) for all $a \in A$, $t_a^\#(X) R t_a^\#(Y)$. Since this is just the old definition of bisimulation (Definition 1) applied to $(\mathcal{P}(S), o^\#, t^\#)$, we get that $X \sim Y$ iff there exists a bisimulation relating them.

Remark 3 (Linear time v.s. branching time). *It is important not to confuse these bisimulation relations with the standard Milner-and-Park bisimulations [24] (which strictly imply language equivalence): in a standard bisimulation R , if the following states x and y of an NFA are in R ,*

then each x_i should be in R with some y_j (and vice-versa). Here, instead, we first transform the transition relation into

$$x \xrightarrow{a} x_1 + \dots + x_n \quad y \xrightarrow{a} y_1 + \dots + y_m,$$

using the powerset construction, and then we require that the sets $x_1 + \dots + x_n$ and $y_1 + \dots + y_m$ are related by R .

3.2 Bisimulation up to congruence

The semi-lattice structure $(\mathcal{P}(S), +, 0)$ carried by determinised NFA makes it possible to introduce a new up-to technique, which is not available with plain DFA: *up to congruence*. This technique relies on the following function.

Definition 5 (Congruence closure). *Let $u: \mathcal{P}(\mathcal{P}(S) \times \mathcal{P}(S)) \rightarrow \mathcal{P}(\mathcal{P}(S) \times \mathcal{P}(S))$ be the function on relations on sets of states defined for all $R \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ as:*

$$u(R) = \{(X_1 + X_2, Y_1 + Y_2) \mid X_1 R Y_1 \text{ and } X_2 R Y_2\}.$$

The function $c = (r \cup s \cup t \cup u \cup \text{id})^\omega$ is called the congruence closure function.

Intuitively, $c(R)$ is the smallest equivalence relation which is closed with respect to $+$ and which includes R . It could alternatively be defined inductively using the rules r , s , t , and id from the previous section, and the following one:

$$\frac{X_1 c(R) Y_1 \quad X_2 c(R) Y_2}{X_1 + X_2 c(R) Y_1 + Y_2} u$$

We call bisimulations up to congruence the bisimulations up to c . We report the explicit definition for the sake of clarity:

Definition 6 (Bisimulation up to congruence). *A bisimulation up to congruence for a NFA (S, o, t) is a relation $R \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ on sets of states, such that whenever $X R Y$ then*

1. $o^\#(X) = o^\#(Y)$ and
2. for all $a \in A$, $t_a^\#(X) c(R) t_a^\#(Y)$.

We then show that bisimulations up to congruence are sound, using the notion of compatibility:

Lemma 3. *The function u is compatible.*

Proof. We assume that $R \mapsto R'$, and we prove that $u(R) \mapsto u(R')$. If $X u(R) Y$, then $X = X_1 + X_2$ and $Y = Y_1 + Y_2$ for some X_1, X_2, Y_1, Y_2 such that $X_1 R Y_1$ and $X_2 R Y_2$. By assumption, we have $o^\#(X_1) = o^\#(Y_1)$, $o^\#(X_2) = o^\#(Y_2)$, and for all $a \in A$, $t_a^\#(X_1) R' t_a^\#(Y_1)$ and $t_a^\#(X_2) R' t_a^\#(Y_2)$. Since $o^\#$ and $t^\#$ are homomorphisms, we deduce $o^\#(X_1 + X_2) = o^\#(Y_1 + Y_2)$, and for all $a \in A$, $t_a^\#(X_1 + X_2) u(R') t_a^\#(Y_1 + Y_2)$. $\square$

Theorem 2. *Any bisimulation up to congruence is contained in a bisimulation.*

Proof. By Proposition 3, it suffices to show that c is compatible, which follows from Lemmas 1, 2 and 3. $\square$

Figure 3. Bisimulations up to congruence, on a single letter NFA.

In the Introduction, we already gave an example of bisimulation up to context, which is a particular case of bisimulation up to congruence (up to context corresponds to use just the function $(r \cup u \cup \text{id})^\omega$, without closing under s and t).

A more involved example illustrating the use of all ingredients of the congruence closure function (c) is given in Figure 3. The relation R expressed by the dashed numbered lines (formally $R = \{(x, u), (y + z, u)\}$) is neither a bisimulation, nor a bisimulation up to equivalence, since $y + z \xrightarrow{a} x + y$ and $u \xrightarrow{a} u$, but $(x + y, u) \notin c(R)$. However, R is a bisimulation up to congruence. Indeed, we have $(x + y, u) \in c(R)$:

$$\begin{aligned} x + y & c(R) u + y & ((x, u) \in R) \\ c(R) y + z + y & & ((y + z, u) \in R) \\ & = y + z \\ c(R) u & & ((y + z, u) \in R) \end{aligned}$$

In contrast, we need four pairs to get a bisimulation up to c containing (x, u) : this is the relation depicted with both dashed and dotted lines in Figure 3.

Note that we can deduce many other equations from R ; in fact, $c(R)$ defines the following partition of sets of states:

$$\{0\}, \{y\}, \{z\}, \{x, u, x + y, x + z, \text{ and the 9 remaining subsets}\}.$$

3.3 Optimised algorithm for NFA

Algorithms for NFA can be obtained by computing the determinised NFA on-the-fly [12]: starting from the algorithms for DFA (Figure 1), it suffices to work with sets of states, and to inline the powerset construction. The corresponding code is given in Figure 4. The naive algorithm (Naive) does not use any up to technique, Hopcroft and Karp's algorithm (HK) reasons up to equivalence in step 3.2, and the optimised algorithm, referred as HKC in the sequel, relies on up to congruence: step 3.2 becomes

(3.2) **if** $(X', Y') \in c(R \cup \text{todo})$ **then skip**;

Observe that we use $c(R \cup \text{todo})$ rather than $c(R)$: this allows us to skip more pairs, and this is safe since all pairs in todo will eventually be processed.

Corollary 2. *For all $X, Y \in \mathcal{P}(S)$, $X \sim Y$ iff HKC(X, Y).*

Proof. Same proof as for Proposition 2, by using the invariant $R \mapsto c(R \cup \text{todo})$ for the loop. We deduce that R is a bisimulation up to congruence after the loop. We conclude with Theorem 2 and Proposition 1. $\square$

The most important point about these three algorithms is that they compute the states of the determinised NFA lazily. This means that only *accessible* states need to be computed, which is of practical importance since the determinised NFA can be exponentially large. In case of a negative answer, the three algorithms stop even before all accessible states have been explored; otherwise, if a bisimulation (possibly up-to) is found, it depends on the algorithm:

Naive(X, Y)

```

(1)  $R$  is empty;  $todo$  is empty;
(2) insert  $(X, Y)$  in  $todo$ ;
(3) while  $todo$  is not empty, do {
  (3.1) extract  $(X', Y')$  from  $todo$ ;
  (3.2) if  $(X', Y') \in R$  then skip;
  (3.3) if  $o^\sharp(X') \neq o^\sharp(Y')$  then return false;
  (3.4) for all  $a \in A$ ,
    insert  $(t_a^\sharp(X'), t_a^\sharp(Y'))$  in  $todo$ ;
  (3.5) insert  $(X', Y')$  in  $R$ ;
(4) return true;

```

Figure 4. On-the-fly naive algorithm, for checking the equivalence of sets of states X and Y of a NFA (S, o, t) . The code for on-the-fly HK(X, Y) is obtained by replacing the test in step 3.2 with $(X', Y') \in e(R)$; the code for HKC(X, Y) is obtained by replacing this test with $(X', Y') \in c(R \cup todo)$.

- with Naive, all accessible states need to be visited, by definition of bisimulation;
- with HK, the only case where some accessible states can be avoided is when a pair (X, X) is encountered: the algorithm skips this pair so that the successors of X are not necessarily computed (this situation rarely happens in practice—it actually never happens when starting with disjoint automata). In the other cases where a pair (X, Y) is skipped, then X and Y are necessarily already related to some other states in R , so that their successors will eventually be explored;
- with HKC, only a small portion of the accessible states is built (check the experiments in Section 6). To see a concrete example, let us execute HKC on the NFA from Figure 3. After two iterations, $R = \{(x, u), (y + z, u)\}$. Since $x + y \not\sim c(R) u$, the algorithm stops without building the states $x + y$ and $x + y + z$. Similarly, in the example from the Introduction, HKC does not construct the four states corresponding to pairs 4, 5, and 6.

This ability of HKC to ignore parts of the determinised NFA comes from the up to congruence technique, which allows one to infer properties about states that were not necessarily encountered before. As we shall see in Section 4 the efficiency of antichains algorithms [1, 33] also comes from their ability to skip large parts of the determinised NFA.

3.4 Computing the congruence closure

For the optimised algorithm to be effective, we need a way to check whether some pairs belong to the congruence closure of some relation (step 3.2). We present here a simple solution based on set rewriting; the key idea is to look at each pair (X, Y) in a relation R as a pair of rewriting rules:

$$X \rightarrow X + Y \quad Y \rightarrow X + Y,$$

which can be used to compute normal forms for sets of states. Indeed, by idempotence, $X R Y$ entails $X \sim c(R) X + Y$.

Definition 7. Let $R \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ be a relation on sets of states. We define $\sim_R \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ as the smallest irreflexive relation that satisfies the following rules:

$$\frac{X R Y}{X \sim_R X + Y} \quad \frac{X R Y}{Y \sim_R X + Y} \quad \frac{Z \sim_R Z'}{U + Z \sim_R U + Z'}$$

Lemma 4. For all relations R , the relation $\sim_R$ is convergent.

In the sequel, we denote by $X \downarrow_R$ the normal form of a set X w.r.t. $\sim_R$. Intuitively, the normal form of a set is the largest set

of its equivalence class. Recalling the example from Figure 3, the common normal form of $x + y$ and u can be computed as follows (R is the relation $\{(x, u), (y + z, u)\}$):

$$\begin{array}{c} x + y \quad \quad \quad x + y + u \quad \quad \quad x + u \\ \quad \quad \quad \swarrow \quad \quad \quad \swarrow \quad \quad \quad \swarrow \\ \quad \quad \quad x + y + z + u \end{array}$$

Theorem 3. For all relations R , and for all $X, Y \in \mathcal{P}(S)$, we have $X \downarrow_R = Y \downarrow_R$ iff $(X, Y) \in c(R)$.

Thus, in order to check if $(X, Y) \in c(R \cup todo)$ we only have to compute the normal form of X and Y with respect to $\sim_{R \cup todo}$. Note that each pair of $R \cup todo$ may be used only once as a rewriting rule, but we do not know in advance in which order to apply these rules. Therefore, the time required to find one rule that applies is in the worst case rn where $r = |R \cup todo|$ is the size of the relation $R \cup todo$, and $n = |S|$ is the number of states of the NFA (assuming linear time complexity for set-theoretic union and containment of sets of states). Since we cannot apply more than r rules, the time for checking whether $(X, Y) \in c(R \cup todo)$ is bounded by $r^2 n$.

We tried other solutions, notably by using binary decision diagrams [8]. We have chosen to keep the presented rewriting algorithm for its simplicity and because it behaves well in practice.

3.5 Complexity hints

The complexity of Naive, HK and HKC is closely related to the size of the relation that they build. Hereafter, we use $v = |A|$ to denote the number of letters in A .

Lemma 5. The three algorithms require at most $1 + v \cdot |R|$ iterations, where $|R|$ is the size of the produced relation; moreover, this bound is reached whenever they return true.

Therefore, we can conveniently reason about $|R|$.

Lemma 6. Let R_{Naive} , R_{HK} , and R_{HKC} denote the relations produced by the three algorithms. We have

$$|R_{HKC}|, |R_{HK}| \leq m \quad |R_{Naive}| \leq m^2, \quad (2)$$

where $m \leq 2^n$ is the number of accessible states in the determinised NFA and n is the number of states of the NFA. If the algorithms returned true, we moreover have

$$|R_{HKC}| \leq |R_{HK}| \leq |R_{Naive}|. \quad (3)$$

As shown below in Section 4.2.4, R_{HKC} can be exponentially smaller than R_{HK} . Notice however that the problem of deciding NFA language equivalence is PSPACE-complete [23], and that none of the algorithms presented here is in PSPACE: all of them store a set of visited pairs, and in the worst case, this set can become exponentially large with all of them. (This also holds for the antichain algorithms [1, 33] which we describe in Section 4.) Instead, the standard PSPACE algorithm does not store any set of visited pairs: it checks all words of length smaller than 2^n . While this can be done in polynomial space, this systematically requires exponential time.

3.6 Using HKC for checking language inclusion

For NFA, language inclusion can be reduced to language equivalence in a rather simple way. Since the function $\llbracket - \rrbracket : \mathcal{P}(S) \rightarrow 2^{A^*}$ is a semi-lattice homomorphism (see Theorem 7 in Appendix A), for any given sets of states X and Y , $\llbracket X + Y \rrbracket = \llbracket Y \rrbracket$ iff $\llbracket X \rrbracket + \llbracket Y \rrbracket = \llbracket Y \rrbracket$ iff $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$. Therefore, it suffices to run HKC($X + Y, Y$) to check the inclusion $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$.

In such a situation, all pairs that are eventually manipulated by HKC have the shape $(X' + Y', Y')$ for some sets X', Y' . The step 3.2 of HKC, where it checks whether the current pair belongs

to the congruence closure of the relation, can thus be simplified. First, the pairs in the current relation can only be used to rewrite from right to left. Second, the following lemma allows one to avoid unnecessary normal form computations:

Lemma 7. *For all sets X, Y and for all relations R , we have $X + Y \sqsubseteq c(R) Y$ iff $X \subseteq Y \downarrow_R$.*

Proof. We first prove that for all X, Y , $X \downarrow_R = Y \downarrow_R$ iff $X \subseteq Y \downarrow_R$ and $Y \subseteq X \downarrow_R$, using the fact that the normalisation function $\downarrow_R: X \mapsto X \downarrow_R$ is monotone and idempotent. The announced result follows by Theorem 3, since $Y \subseteq (X + Y) \downarrow_R$ is always true and $X + Y \subseteq Y \downarrow_R$ iff $X \subseteq Y \downarrow_R$. $\square$

However, as shown below, checking an equivalence by decomposing it into two inclusions cannot be more efficient than checking the equivalence directly.

Lemma 8. *Let X, Y be two sets of states; let $R_{\subseteq}$ and $R_{\supseteq}$ be the relations computed by $\text{HKC}(X + Y, Y)$ and $\text{HKC}(X + Y, X)$, respectively. If $R_{\subseteq}$ and $R_{\supseteq}$ are bisimulations up to congruence, then the following relation is a bisimulation up to congruence:*

$$R = \{(X', Y') \mid (X' + Y', Y') \in R_{\subseteq} \text{ or } (X' + Y', X') \in R_{\supseteq}\}.$$

On the contrary, checking the equivalence directly actually allows one to skip some pairs that cannot be skipped when reasoning by double inclusion. As an example, consider the DFA on the right of Figure 2. The relation computed by $\text{HKC}(x, u)$ contains only four pairs (because the fifth one follows from transitivity). Instead, the relations built by $\text{HKC}(x, x + u)$ and $\text{HKC}(u + x, u)$ would both contain five pairs: transitivity cannot be used since our relations are now oriented (from $y \leq v$, $z \leq v$ and $z \leq w$, we cannot deduce $y \leq w$). Another example, where we get an exponential factor by checking the equivalence directly rather than through the two inclusions, can be found in Section 4.2.4.

In a sense, the behaviour of the coinduction proof method here is similar to that of standard proofs by induction, where one often has to strengthen the induction predicate to get a (nicer) proof.

4. Antichain algorithm

In [33], De Wulf et al. have proposed the *antichain* approach for checking language inclusion of NFA. We show that this approach can be explained in terms of *simulations up to upward-closure* that, in turn, can be seen as a special case of bisimulations up to congruence. Before doing so, we recall the standard notion of antichain and we describe the antichain algorithm (AC).

Given a partial order $(X, \sqsubseteq)$, an *antichain* is a subset $Y \subseteq X$ containing only incomparable elements (that is, for all $y_1, y_2 \in Y$, $y_1 \not\sqsubseteq y_2$ and $y_2 \not\sqsubseteq y_1$). AC exploits antichains over the set $S \times \mathcal{P}(S)$, where the ordering is given by $(x_1, Y_1) \sqsubseteq (x_2, Y_2)$ iff $x_1 = x_2$ and $Y_1 \subseteq Y_2$.

In order to check $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ for two sets of states X, Y of an NFA (S, o, t) , AC maintains an antichain of pairs (x', Y') , where x' is a state of the NFA and Y' is a state of the determinised automaton. More precisely, the automaton is explored non-deterministically (via t) for obtaining the first component of the pair and deterministically (via $t^\#$) for the second one. If a pair such that x' is accepting ($o(x') = 1$) and Y' is not ($o^\#(Y') = 0$) is encountered, then a counter-example has been found. Otherwise all derivatives of the pair along the automata transitions have to be inserted into the antichain, so that they will be explored. If one these pairs p is larger than a previously encountered pair p' ($p' \sqsubseteq p$) then the language inclusion corresponding to p is subsumed by p' so that p can be skipped; otherwise, if $p \sqsubseteq p_1, \dots, p_n$ for some pairs

$p_1, \dots, p_n$ that are already in the antichain, then one can safely remove these pairs: they are subsumed by p and, by doing so, the set of visited pairs remains an antichain.

Remark 4. *An important difference between HKC and AC consists in the fact that the former inserts pairs in *todo* without checking whether they are redundant (this check is performed when the pair is processed), while the latter removes all redundant pairs whenever a new one is inserted. Therefore, the cost of an iteration with HKC is merely the cost of the corresponding congruence check, while the cost of an iteration with AC is merely that of inserting all successors of the corresponding pair and simplifying the antichain.*

Note that the above description corresponds to the “forward” antichain algorithm, as described in [1]. Instead, the original antichain algorithm, as first described in [33], is “backward” in the sense that the automata are traversed in the reversed way, from accepting states to initial states. The two versions are dual [33] and we could similarly define the backward counterpart of HKC and HK. We however stick to the forward versions for the sake of clarity.

4.1 Coinductive presentation

Leaving apart the concrete data structures used to manipulate antichains, we can rephrase this algorithm using a coinductive framework, like we did for Hopcroft and Karp’s algorithm.

First define a notion of *simulation*, where the left-hand side automaton is executed non-deterministically:

Definition 8 (Simulation). *Given two relations $T, T' \subseteq S \times \mathcal{P}(S)$, T s-progresses to T' , denoted $T \rightsquigarrow_s T'$, if whenever $x T Y$ then*

1. $o(x) \leq o^\#(Y)$ and
2. for all $a \in A$, $x' \in t_a(x)$, $x' T' t_a^\#(Y)$.

A simulation is a relation T such that $T \rightsquigarrow_s T$.

As expected, we obtain the following coinductive proof principle:

Proposition 4 (Coinduction). *For all sets X, Y , we have $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ iff there exists a simulation T such that for all $x \in X$, $x T Y$.*

(Note that like for our notion of bisimulation, the above notion of simulation is weaker than the standard one from concurrency theory [24], which *strictly* entails language inclusion—Remark 3.)

To account for the antichain algorithm, where we can discard pairs using the preorder $\sqsubseteq$, it suffices to define the *upward closure* function $\uparrow: \mathcal{P}(S \times \mathcal{P}(S)) \rightarrow \mathcal{P}(S \times \mathcal{P}(S))$ as

$$\uparrow T = \{(x, Y) \mid \exists (x', Y') \in T \text{ s.t. } (x', Y') \sqsubseteq (x, Y)\}.$$

A pair belongs to the upward closure $\uparrow T$ of a relation $T \subseteq S \times \mathcal{P}(S)$, if and only if this pair is subsumed by some pair in T . In fact, rather than trying to construct a simulation, AC attempts to construct a simulation up to upward closure.

Like for HK and HKC, this method can be justified by defining the appropriate notion of s-compatible function, showing that any simulation up to an s-compatible function is contained in a simulation, and showing that the upward closure function ($\uparrow$) is s-compatible.

Theorem 4. *Any simulation up to $\uparrow$ is contained in a simulation.*

Corollary 3. *For all $X, Y \in \mathcal{P}(S)$, $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ iff $\text{AC}(X, Y)$.*

4.2 Comparing HKC and AC

The efficiency of the two algorithms strongly depends on the number of pairs that they need to explore. In the following (Sections 4.2.3 and 4.2.4), we show that HKC can explore far fewer pairs than AC, when checking language inclusion of automata that share some states, or when checking language equivalence. We would also like to formally prove that (a) HKC never explores more than AC, and

(b) when checking inclusion of disjoint automata, AC never explores more than HKC. Unfortunately, the validity of these statements highly depends on numerous assumptions about the two algorithms (e.g., on the exploration strategy) and their potential proofs seem complicated and not really informative. For these reasons, we preferred to investigate the formal correspondence at the level of the coinductive proof techniques, where it is much cleaner.

4.2.1 Language inclusion: HKC can mimic AC

As explained in Section 3.6, we can check the language inclusion of two sets X, Y by executing $\text{HKC}(X+Y, Y)$. We now show that for any simulation up to upward closure that proves the inclusion $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$, there exists a bisimulation up to congruence of the same size which proves the same inclusion. For $T \subseteq S \times \mathcal{P}(S)$, let $\hat{T} \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ denote the relation $\{(x+Y, Y) \mid x T Y\}$.

Lemma 9. *We have $\hat{\uparrow}T \subseteq c(\hat{T})$.*

Proof. If $(x+Y, Y) \in \hat{\uparrow}T$, then there exists $Y' \subseteq Y$ such that $(x, Y') \in T$. By definition, $(x+Y', Y') \in \hat{T}$ and $(Y, Y) \in c(\hat{T})$. By the rule (u), $(x+Y'+Y, Y'+Y) \in c(\hat{T})$ and since $Y' \subseteq Y$, $(x+Y, Y) \in c(\hat{T})$. $\square$

Proposition 5. *If T is a simulation up to $\uparrow$, then $\hat{T}$ is a bisimulation up to c .*

Proof. First observe that if $T \rightsquigarrow_s T'$, then $\hat{T} \rightsquigarrow u^\omega(\hat{T}')$. Therefore, if $T \rightsquigarrow_s \uparrow T$, then $\hat{T} \rightsquigarrow u^\omega(\hat{\uparrow}T)$. By Lemma 9, $\hat{T} \rightsquigarrow u^\omega(c(\hat{T})) = c(\hat{T})$. $\square$

(Note that transitivity and symmetry are not used in the above proofs: the constructed bisimulation up to congruence is actually a bisimulation up to context $(r \cup u \cup id)^\omega$.)

The relation $\hat{T}$ is not the one computed by HKC, since the former contains pairs of the shape $(x+Y, Y)$, while the latter has pairs of the shape $(X+Y, Y)$ with X possibly not a singleton. However, note that manipulating pairs of the two kinds does not change anything since by Lemma 7, $(X+Y, Y) \in c(R)$ iff for all $x \in X$, $(x+Y, Y) \in c(R)$.

4.2.2 Inclusion: AC can mimic HKC on disjoint automata

As shown in Section 4.2.3 below, HKC can be faster than AC, thanks to the up to transitivity technique. However, in the special case where the two automata are disjoint, transitivity cannot help, and the two algorithms actually match each other.

Suppose that the automaton (S, o, t) is built from two disjoint automata (S_1, o_1, t_1) and (S_2, o_2, t_2) as described in Remark 1. Let R be the relation obtained by running $\text{HKC}(X_0+Y_0, Y_0)$ with $X_0 \subseteq S_1$ and $Y_0 \subseteq S_2$. All pairs in R are necessarily of the shape $(X+Y, Y)$ with $X \subseteq S_1$ and $Y \subseteq S_2$. Let $\bar{R} \subseteq S \times \mathcal{P}(S)$ denote the relation $\{(x, Y) \mid \exists X, x \in X \text{ and } X+Y \bar{R} Y\}$.

Lemma 10. *If S_1 and S_2 are disjoint, then $\bar{c}(\bar{R}) \subseteq \uparrow(\bar{R})$.*

Proof. Suppose that $x \bar{c}(\bar{R}) Y$, i.e., $x \in X$ with $X+Y \bar{c}(\bar{R}) Y$. By Lemma 7, we have $X \subseteq Y \downarrow_R$, and hence, $x \in Y \downarrow_R$. By definition of R the pairs it contains can only be used to rewrite from right to left; moreover, since S_1 and S_2 are disjoint, such rewriting steps cannot enable new rewriting rules, so that all steps can be performed in parallel: we have $Y \downarrow_R = \sum_{X'+Y'RY' \subseteq Y} X'$. Therefore, there exists some X', Y' with $x \in X'$, $X'+Y' \bar{R} Y'$, and $Y' \subseteq Y$. It follows that $(x, Y') \in \bar{R}$, hence $(x, Y) \in \uparrow(\bar{R})$. $\square$

Proposition 6. *If S_1 and S_2 are disjoint, and if R is a bisimulation up to congruence, then $\bar{R}$ is a simulation up to upward closure.*

Figure 5. Family of examples where HKC exponentially improves over AC and HK; we have $x + y \sim z$.

Proof. First observe that for all relations R, R' , if $R \rightsquigarrow R'$, then $\bar{R} \rightsquigarrow_s \bar{R}'$. Therefore, if $R \rightsquigarrow c(R)$, then $\bar{R} \rightsquigarrow_s c(\bar{R})$. We deduce $\bar{R} \rightsquigarrow_s \uparrow(\bar{R})$ by Lemma 10. $\square$

4.2.3 Inclusion: AC cannot mimic HKC on merged automata

The containment of Lemma 10 does not hold when S_1 and S_2 are not disjoint, since c can exploit transitivity, while $\uparrow$ cannot. For a concrete grasp, take $R = \{(x+y, y), (y+z, z)\}$ and observe that $(x, z) \in c(\bar{R})$ but $(x, z) \notin \uparrow(\bar{R})$. This difference makes it possible to find bisimulations up to c that are much smaller than the corresponding simulations up to $\uparrow$, and for HKC to be more efficient than AC. Such an example, where HKC is exponentially better than AC for checking language inclusion of automata sharing some states, is given in [6].

4.2.4 Language equivalence: AC cannot mimic HKC.

AC can be used to check language equivalence, by checking the two underlying inclusions. However, checking equivalence directly can be better, even in the disjoint case. To see this on a simple example, consider the DFA on the right-hand side of Figure 2. If we use AC twice to prove $x \sim u$, we get the following antichains

$$T_1 = \{(x, u), (y, v), (y, w), (z, v), (z, w)\}, \\ T_2 = \{(u, x), (v, y), (w, y), (v, z), (w, z)\},$$

containing five pairs each. Instead, four pairs are sufficient with HK or HKC, thanks to up to symmetry and up to transitivity.

For a more interesting example, consider the family of NFA given in Figure 5, where n is an arbitrary natural number. Taken together, the states x and y are equivalent to the state z : they recognise the language $(a+b)^*(a+b)^{n+1}$. Alone, the state x (resp. y) recognises the language $(a+b)^*a(a+b)^n$ (resp. $(a+b)^*b(a+b)^n$).

For $i \leq n$, let $X_i = x+x_1+\dots+x_i$, $Y_i = y+y_1+\dots+y_i$, and $Z_i = z+z_1+\dots+z_i$; for $N \subseteq [1..n]$, furthermore set

$$X_i^N = x + \sum_{j \in N} x_j, \quad \bar{Y}_i^N = y + \sum_{j \in [1..n] \setminus N} y_j.$$

In the determinised NFA, $x+y$ can reach all the states of the shape $X_i^N + \bar{Y}_i^N$, for $i \leq n$ and $N \subseteq [1..i]$. For instance, for $n=i=2$, we have $x+y \xrightarrow{aa} x+y+x_1+x_2$, $x+y \xrightarrow{ab} x+y+y_1+x_2$, $x+y \xrightarrow{ba} x+y+x_1+y_2$, and $x+y \xrightarrow{bb} x+y+y_1+y_2$. Instead, z reaches only $n+1$ distinct states, those of the form Z_i .

The smallest bisimulation relating $x+y$ and z is

$$R = \{(X_i^N + \bar{Y}_i^N, Z_i) \mid i \leq n, N \subseteq [1..i]\},$$

which contains $2^{n+1}-1$ pairs. This is the relation computed by $\text{Naive}(x, y)$ and $\text{HK}(x, y)$ —the up to equivalence technique (alone) does not help in HK. With AC, we obtain the antichains $T_x + T_y$ (for

$\llbracket x + y \rrbracket \subseteq \llbracket z \rrbracket$) and T_z (for $\llbracket x + y \rrbracket \supseteq \llbracket z \rrbracket$), where:

$$\begin{aligned} T_x &= \{(x_i, Z_i) \mid i \leq n\}, \\ T_y &= \{(y_i, Z_i) \mid i \leq n\}, \\ T_z &= \{(z_i, X_i^N + \bar{Y}_i^N) \mid i \leq n, N \subseteq [1..i]\}. \end{aligned}$$

Note that T_x and T_y have size $n + 1$, and T_z has size $2^{n+1} - 1$.

The language recognised by x or y are known for having a minimal DFA with 2^n states [17]. So, checking $x + y \sim z$ via minimisation (e.g., [9, 15]) would also require exponential time.

This is not the case with HKC, which requires only polynomial time in this case. Indeed, $\text{HKC}(x+y, z)$ builds the relation

$$\begin{aligned} R' &= \{(x + y, z)\} \\ &\cup \{(x + Y_i + y_{i+1}, Z_{i+1}) \mid i < n\} \\ &\cup \{(x + Y_i + x_{i+1}, Z_{i+1}) \mid i < n\} \end{aligned}$$

which is a bisimulation up to congruence and which only contains $2n + 1$ pairs. To see that this is a bisimulation up to congruence, consider the pair $(x+y+x_1+y_2, Z_2)$ obtained from $(x+y, z)$ after reading the word ba . This pair does not belong to R' but to its congruence closure. Indeed, we have

$$\begin{aligned} x+y+x_1+y_2 &c(R') Z_1+y_2 && (x+y+x_1 R' Z_1) \\ &c(R') x+y+y_1+y_2 && (x+y+y_1 R' Z_1) \\ &c(R') Z_2 && (x+y+y_1+y_2 R' Z_2) \end{aligned}$$

(Check Lemma 18 in Appendix D for a complete proof.)

5. Exploiting Similarity

Looking at the example in Figure 5, a natural idea would be to first quotient the automaton by graph isomorphism. By doing so, we would merge the states x_i, y_i, z_i , and we would obtain the following automaton, for which checking $x+y \sim z$ is much easier.

As shown by Abdulla et al. [1], one can actually do better with the antichain algorithm, by exploiting any preorder contained in language inclusion (e.g., similarity [24]). In this section, we rephrase this technique for antichains in our coinductive framework, and we show how this idea can be embedded in HKC, resulting in an even stronger algorithm.

5.1 AC with similarity: AC'

For the sake of clarity, we fix the preorder to be *similarity*, which can be computed in quadratic time [13]:

Definition 9 (Similarity). *Similarity is the largest relation on states $\preceq \subseteq S \times S$ such that $x \preceq y$ entails:*

1. $o(x) \leq o(y)$ and
2. for all $a \in A, x' \in S$ such that $x \xrightarrow{a} x'$, there exists some y' such that $y \xrightarrow{a} y'$ and $x' \preceq y'$.

One extends similarity to a preorder $\preceq^{\forall\exists} \subseteq \mathcal{P}(S) \times \mathcal{P}(S)$ on sets of states, and to a preorder $\sqsubseteq^{\preceq} \subseteq (S \times \mathcal{P}(S)) \times (S \times \mathcal{P}(S))$ on antichain pairs, as:

$$\begin{aligned} X \preceq^{\forall\exists} Y &\text{ if } \forall x \in X, \exists y \in Y, x \preceq y, \\ (x', Y') \sqsubseteq^{\preceq} (x, Y) &\text{ if } x \preceq x' \text{ and } Y' \preceq^{\forall\exists} Y. \end{aligned}$$

The new antichain algorithm [1], which we call AC' , is similar to AC , but the antichain is now taken w.r.t. the new preorder $\sqsubseteq^{\preceq}$. Formally, let $\hat{\lambda}: \mathcal{P}(S \times \mathcal{P}(S)) \rightarrow \mathcal{P}(S \times \mathcal{P}(S))$ be the function defined for all relations $T \subseteq S \times \mathcal{P}(S)$, as

$$\begin{aligned} \hat{\lambda}T &= \{(x, Y) \mid x \preceq^{\forall\exists} Y, \text{ or} \\ &\quad \exists (x', Y') \in T \text{ s.t. } (x', Y') \sqsubseteq^{\preceq} (x, Y)\}. \end{aligned}$$

While AC consists in trying to build a simulation up to $\hat{\lambda}$, AC' tries to build a simulation up to $\hat{\lambda}$, i.e., it skips a pair (x, Y) if either (a) it is subsumed by another pair of the antichain or (b) $x \preceq^{\forall\exists} Y$.

Theorem 5. *Any simulation up to $\hat{\lambda}$ is contained in a simulation.*

Corollary 4. *The antichain algorithm proposed in [1] is sound and complete: for all sets $X, Y, \llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ iff $\text{AC}'(X, Y)$.*

Optimisation 1(a) and optimisation 1(b) in [1] are simply (a) and (b), as discussed above. Another optimisation, called Optimisation 2, is presented in [1]: if $y_1 \preceq y_2$ and $y_1, y_2 \in Y$ for some pair (x, Y) , then y_1 can be safely removed from Y . Note that while this is useful to store smaller sets, it does not allow one to explore less, since the pairs encountered with or without optimisation 2 are always equivalent w.r.t. the ordering $\sqsubseteq^{\preceq}$: $Y \preceq^{\forall\exists} Y \setminus y_1$ and, for all $a \in A, t_a^\#(Y) \preceq^{\forall\exists} t_a^\#(Y \setminus y_1)$.

5.2 HKC with similarity: HKC'

Although HKC is primarily designed to check language equivalence, we can also extend it to exploit the similarity preorder. It suffices to notice that for any similarity pair $x \preceq y$, we have $x+y \sim y$.

Let $\overline{\preceq}$ denote the relation $\{(x+y, y) \mid x \preceq y\}$, let r' denote the constant to $\overline{\preceq}$ function, and let $c' = (r' \cup s \cup t \cup u \cup \text{id})^\omega$. Accordingly, we call HKC' the algorithm obtained from HKC (Figure 4) by replacing $(X, Y) \in c(R \cup \text{todo})$ with $(X, Y) \in c'(R \cup \text{todo})$ in step 3.2. Notice that the latter test can be reduced to rewriting thanks to Theorem 3 and the following lemma.

Lemma 11. *For all relations $R, c'(R) = c(R \cup \overline{\preceq})$.*

In other words to check whether $(X, Y) \in c'(R \cup \text{todo})$, it suffices to compute the normal forms of X and Y w.r.t. the rules from $R \cup \text{todo}$ plus the rules $x + y \leftarrow y$ for all $x \preceq y$.

Theorem 6. *Any bisimulation up to c' is contained in a bisimulation.*

Proof. Consider the constant function $r'': \mathcal{P}(\mathcal{P}(S) \times \mathcal{P}(S)) \rightarrow \mathcal{P}(\mathcal{P}(S) \times \mathcal{P}(S))$ mapping all relations to $\sim$. Since language equivalence ($\sim$) is a bisimulation, we immediately obtain that this function is compatible. Thus so is the function $c'' = (r'' \cup s \cup t \cup u \cup \text{id})^\omega$. We have that $\overline{\preceq}$ is contained in $\sim$, so that any bisimulation up to c' is a bisimulation up to c'' . Since c'' is compatible, such a relation is contained in a bisimulation, by Proposition 3. $\square$

Note that in the above proof, we can replace $\overline{\preceq}$ by any other relation contained in $\sim$. Intuitively, bisimulations up to c'' correspond to classical bisimulations up to bisimilarity [24] from concurrency.

Corollary 5. *For all sets X, Y , we have $X \sim Y$ iff $\text{HKC}'(X, Y)$.*

5.3 Relationship between HKC' and AC'

Like in Section 4.2.1, we can show that for any simulation up to $\hat{\lambda}$ there exists a corresponding bisimulation up to c' , of the same size.

Lemma 12. *For all relations $T \subseteq S \times \mathcal{P}(S), \hat{\lambda}\hat{T} \subseteq c'(\hat{T})$.*

Proposition 7. *If T is a simulation up to $\hat{\lambda}$, then $\hat{T}$ is a bisimulation up to c' .*

Figure 6. Family of examples where HKC' exponentially improves over AC' , for inclusion of disjoint automata: we have $\llbracket z \rrbracket \subseteq \llbracket x+y \rrbracket$.

However, even for checking inclusion of disjoint automata, AC' cannot mimic HKC' , because now the similarity relation allows one to exploit transitivity. To see this, consider the example given in Figure 6, where we want to check that $\llbracket z \rrbracket \subseteq \llbracket x+y \rrbracket$, and for which the similarity relation is shown on the right-hand side.

Since this is an inclusion of disjoint automata, HKC and AC , which do not exploit similarity, behave the same (cf. Sections 4.2.1 and 4.2.2). Actually, they also behave like HK and they require $2^{n+1} - 1$ pairs. On the contrary, the use of similarity allows HKC' to prove the inclusion with only $2n + 1$ pairs, by computing the following bisimulation up to c' (Lemma 19 in Appendix E):

$$\begin{aligned} R'' = & \{(z+x+y, x+y)\} \\ & \cup \{(Z_{i+1}+X_i+y+y_{i+1}, X_i+y+y_{i+1}) \mid i < n\} \\ & \cup \{(Z_{i+1}+X_{i+1}+y, X_{i+1}+y) \mid i < n\}, \end{aligned}$$

where $X_i = x+x_1+\dots+x_i$ and $Z_i = z+z_1+\dots+z_i$.

Like in Section 4.2.4, to see that this is a bisimulation up to c' (where we do exploit similarity), consider the pair obtained after reading the word ab : $(Z_2+x+y+x_2+y_1, x+y+x_2+y_1)$. This pair does not belong to R'' or $c(R'')$, but it does belong to $c'(R'')$. Indeed, by Lemmas 7 and 11, this pair belongs to $c'(R'')$ iff $Z_2 \subseteq (x+y+x_2+y_1) \downarrow_{R'' \cup \subseteq}$, and we have

$$\begin{aligned} & x+y+x_2+y_1 \\ \rightsquigarrow_{R'' \cup \subseteq} & Z_1+x+y+y_1+x_2 \quad (Z_1+x+y+y_1 \ R'' \ x+y+y_1) \\ \rightsquigarrow_{R'' \cup \subseteq} & Z_1+X_1+y+y_1+x_2 = Z_1+X_2+y+y_1 \quad (x_1 \preceq z_1) \\ \rightsquigarrow_{R'' \cup \subseteq} & Z_2+X_2+y+y_1+x_2 \quad (Z_2+X_2+y \ R'' \ X_2+y) \end{aligned}$$

On the contrary, AC' is not able to exploit similarity in this case, and it behaves like AC : both of them compute the same antichain T_z as in the example from Section 4.2.4, which has $2^{n+1} - 1$ elements.

In fact, even when considering inclusion of disjoint automata, the use of similarity tends to virtually merge states, so that HKC' can use the up to transitivity technique which AC and AC' lack.

5.4 A short recap

Figure 7 summarises the relationship amongst the presented algorithms, in the general case and in the special case of language inclusion of disjoint automata. In this diagram, an arrow $X \rightarrow Y$ (from an algorithm X to Y) means that (a) Y can explore less states than X , and (b) Y can mimic X , i.e., the proof technique of Y is at least as powerful as the one of X . (The labels on the arrows point to the sections showing these relations; unlabelled arrows are not illustrated in this paper, they are easily inferred from what we have shown.)

6. Experimental assessment

To get an intuition of the average behaviour of HKC on various NFA, and to compare it with HK and AC , we provide some benchmarks on random automata and on automata obtained from model-checking problems. In both cases, we conduct the experiments on a MacBook pro 2.4GHz Intel Core i7, with 4GB of memory, running OS X

Figure 7. Relationship between the various algorithms.

Lion (10.7.4). We use our OCaml implementation for HK , HKC , and HKC' [6], and the `libvata` C++ library for AC and AC' [20]. (To our knowledge, `libvata` is the most efficient implementation currently available for the antichain algorithms.)

6.1 Random automata

For a given size n , we generate a thousand random NFA with n states and two letters. According to [31], we use a linear transition density of 1.25 (which means that the expected out-degree of each state and with respect to each letter is 1.25): Tabakov and Vardi empirically showed that one statistically gets more challenging NFA with this particular value. We generate NFA without accepting states: by doing so, we make sure that the algorithms never encounter a counter-example, so that they always continue until they find a (bi)simulation up to: these runs correspond to their worst cases for all possible choices of accepting states for the given NFA.²

We run all algorithms on these NFA, starting from two distinct singleton sets, to measure the required time and the number of processed pairs: for HK , HKC , and HKC' , this is the number of pairs put into the bisimulation up to (R) ; for AC and AC' , this is the number of pairs inserted into the antichain. The timings for HKC' and AC' do not include the time required to compute similarity.

We report the median values (50%), the last deciles (90%), the last percentiles (99%), and the maximum values (100%) in Table 1. For instance, for $n = 70$, 90% of the examples require less than 155ms with HK ; equivalently, 10% of the examples require more than 155ms. (For a few tests, `libvata` ran out of memory, whence the ∞ symbols in the table.) We also plotted on Figure 8 the distribution of the number of processed pairs when $n = 100$.

HKC and AC are several orders of magnitude better than HK , and HKC is usually two to ten times faster than AC . Moreover, for the first four lines, HKC is much more predictable than AC , i.e., the last percentiles and maximal values are of the same order as the median value. (AC seems to become more predictable for larger values of n .) The same relative behaviour can be observed between HKC' and AC' ; moreover, HKC alone is apparently faster than AC' .

Also recall that the size of the relations generated by HK is a lower bound for the number of accessible states of the determined NFA (Lemma 6 (2)); one can thus see in Table 1 that HKC usually explores an extremely small portion of these DFA (e.g., less than one per thousand for $n = 100$). The last column reports the median size of the minimal DFA for the corresponding parameters, as given in [31]. HK usually explores much many states than what would be necessary with a minimal DFA, while HKC and AC need much less.

6.2 Automata from model-checking

Checking language inclusion of NFA can be useful for model-checking, where one sometimes has to compute a sequence of NFA

²To get this behaviour for AC and AC' , we actually had to trick `libvata`, which otherwise starts by removing non-coaccessible states, and thus reduces any of these NFA to the empty one.

$n = S $	algo.	required time (seconds)				number of processed pairs				mDFA size
		50%	90%	99%	100%	50%	90%	99%	100%	50%
50	HK	0.007	0.022	0.050	0.119	2511	6299	12506	25272	~1000
	AC	0.002	0.003	0.142	1.083	112	245	2130	5208	
	HKC	0.000	0.000	0.000	0.000	21	26	32	63	
	AC'	0.002	0.002	0.038	0.211	79	131	1098	1926	
	HKC'	0.000	0.000	0.000	0.000	18	23	28	58	
70	HK	0.047	0.155	0.413	0.740	10479	28186	58782	87055	~6000
	AC	0.002	0.003	1.492	4.163	150	285	8383	15575	
	HKC	0.000	0.000	0.000	0.000	27	34	40	49	
	AC'	0.002	0.003	0.320	0.884	110	172	3017	6096	
	HKC'	0.000	0.000	0.000	0.000	23	29	36	44	
100	HK	0.373	1.207	3.435	5.660	58454	164857	361227	471727	~30000
	AC	0.003	0.004	3.214	36.990	204	298	13801	48059	
	HKC	0.000	0.000	0.000	0.001	36	44	54	70	
	AC'	0.003	0.004	0.738	6.966	152	211	4087	18455	
	HKC'	0.000	0.000	0.000	0.001	31	39	46	64	
300	AC	0.009	0.010	0.028	0.750	562	622	2232	14655	—
	HKC	0.001	0.002	0.003	0.009	86	104	118	132	
	AC'	0.012	0.013	0.022	0.970	433	484	920	14160	
	HKC'	0.001	0.001	0.002	0.006	76	91	104	116	
500	AC	0.014	0.015	0.039	∞	918	986	2571	∞	—
	HKC	0.002	0.005	0.008	0.018	130	154	176	193	
	AC'	0.025	0.028	0.042	∞	710	772	1182	∞	
	HKC'	0.002	0.004	0.007	0.013	115	136	154	169	
1000	AC	0.029	0.031	0.038	∞	1808	1878	2282	∞	—
	HKC	0.007	0.022	0.055	0.093	228	271	304	337	
	AC'	0.074	0.080	0.092	∞	1409	1488	1647	∞	
	HKC'	0.008	0.019	0.041	0.077	202	238	265	299	

Table 1. Running the five presented algorithms to check language equivalence on random NFA with two letters.

Figure 8. Distributions of the number of processed pairs, for the 1000 NFA with 100 states and 2 letters from Table 1.

by iteratively applying a transducer, until a fixpoint is reached [7]. To know that the fixpoint is reached, one typically has to check whether an NFA is contained in another one.

Abdulla et al. [1] use such benchmarks to test their algorithm (AC') against the plain antichain algorithm (AC [33]). We reuse them to test HKC' against AC' in a concrete scenario. We take the sequences of automata kindly provided by L. Holik, which roughly corresponds to those used in [1] and which come from the model checking of various programs (the bakery algorithm, bubble sort, and a producer-consumer system). For all these sequences, we check the inclusions of consecutive pairs, in both directions. We separate the results into those for which a counter-example is found, and those for which the inclusion holds. We skip the trivial inclusions which hold by similarity ($\preceq^{\forall\exists}$), and for which both HKC' and AC' stop immediately.

The results are given in Table 2. Even though these are inclusions of disjoint automata, HKC' is faster than AC' on these examples: up to transitivity can be exploited thanks to the similarity pairs, and larger parts of the determinised NFA can be skipped.

7. Related work

A similar notion of bisimulation up to congruence has already been used to obtain decidability and complexity results about context-free processes, under the name of *self-bisimulations*. Caucal [10] introduced this concept to give a shorter and nicer proof of the result by Baeten et al. [4]: bisimilarity is decidable for normed context-free processes. Christensen et al [11] then generalised the result to all context-free processes, also by using self-bisimulations. Hirshfeld et al. [14] used a refinement of this notion to get a polynomial algorithm for bisimilarity in the normed case.

There are two main differences with the ideas we presented here. First, the above papers focus on bisimilarity rather than language equivalence (recall that although we use bisimulation relations, we check language equivalence since we work on the determinised NFA—Remark 3). Second, we consider a notion of bisimulation up to congruence where the congruence is taken with respect to non-determinism (union of sets of states). Self-bisimulations are also bisimulations up to congruence, but the congruence is taken with respect to word concatenation. We cannot consider this operation in our setting since we do not have the corresponding monoid structure in plain NFA.

Other approaches, that are independent from the algebraic structure (e.g., monoids or semi-lattices) and the behavioural equivalence (e.g., bisimilarity or language equivalence) are shown in [5, 21, 22, 26]. These propose very general frameworks into which our up to congruence technique fits as a very special case. To our knowledge, bisimulation up to congruence has never been proposed as a technique for proving language equivalence of NFA.

result	algo.	required time (seconds)				number of processed pairs				number of tests
		50%	90%	99%	100%	50%	90%	99%	100%	
counter-example	AC'	0.012	0.107	1.047	1.134	23	247	598	1352	518
	HKC'	0.001	0.005	0.025	0.383	11	24	112	290	
inclusion holds	AC'	0.079	0.795	1.457	1.480	149	733	1854	3087	178
	HKC'	0.015	0.165	0.340	0.345	61	695	1076	1076	

Table 2. Running HKC' and AC' to test language inclusion of disjoint NFA generated from model-checking.

8. Conclusions and future work

We showed that the standard algorithm by Hopcroft and Karp for checking language equivalence of DFA relies on a bisimulation up to equivalence proof technique; this allowed us to design a new algorithm (HKC) for the non-deterministic case, where we exploit a novel technique called up to congruence.

We then compared HKC to the recently introduced antichain algorithms [33] (AC): when checking the inclusion of disjoint automata, the two algorithms are equivalent, in all the other cases HKC is more efficient since it can use transitivity to prune a larger portion of the state-space.

The difference between these two approaches becomes even more striking when considering some optimisation exploiting similarity. Indeed, as nicely shown with AC' [1], the antichains approach can widely benefit from the knowledge one gets by first computing similarity. Inspired by this work, we showed that both our proof technique (bisimulation up to congruence) and our algorithm (HKC) can be easily modified to exploit similarity. The resulting algorithm (HKC') is now more efficient than AC' even for checking language inclusion of disjoint automata.

We provided concrete examples where HKC and HKC' are exponentially faster than AC and AC' (Sections 4.2.4 and 5.3) and we proved that the coinductive techniques underlying the formers are at least as powerful as those exploited by the latters (Propositions 5 and 7). We finally compared the algorithms experimentally, by running them on both randomly generated automata, and automata resulting from model checking problems. It appears that for these examples, HKC and HKC' perform better than AC and AC'.

Finally note that our implementation of the presented algorithms is available online [6], together with an applet making it possible to test them on user-provided examples.

As future work, we plan to extend our approach to tree automata. In particular, it seems promising to investigate if further up-to techniques can be defined for regular tree expressions. For instance, the algorithms proposed in [3, 18] exploit some optimisation which suggest us coinductive up-to techniques.

References

- [1] P. A. Abdulla, Y.-F. Chen, L. Holík, R. Mayr, and T. Vojnar. When simulation meets antichains. In *Proc. TACAS*, vol. 6015 of *LNCS*, pages 158–174. Springer, 2010.
- [2] A. V. Aho, J. E. Hopcroft, and J. D. Ullman. *The Design and Analysis of Computer Algorithms*. Addison-Wesley, 1974.
- [3] A. Aiken and B. R. Murphy. Implementing regular tree expressions. In *FPCA*, vol. 523 of *LNCS*, pages 427–447. Springer, 1991.
- [4] J. C. M. Baeten, J. A. Bergstra, and J. W. Klop. Decidability of bisimulation equivalence for processes generating context-free languages. In *Proc. PARLE (II)*, vol. 259 of *LNCS*, pages 94–111. Springer, 1987.
- [5] F. Bartels. *On generalized coinduction and probabilistic specification formats*. PhD thesis, Vrije Universiteit Amsterdam, 2004.
- [6] F. Bonchi and D. Pous. Web appendix for this paper. <http://perso.ens-lyon.fr/damien.pous/hknt>, 2012.
- [7] A. Bouajjani, P. Habermehl, and T. Vojnar. Abstract regular model checking. In *Proc. CAV*, vol. 3114 of *LNCS*. Springer, 2004.
- [8] R. E. Bryant. Graph-based algorithms for boolean function manipulation. *IEEE Trans. Computers*, 35(8):677–691, 1986.
- [9] J. A. Brzozowski. Canonical regular expressions and minimal state graphs for definite events. In *Mathematical Theory of Automata*, vol. 12(6), pages 529–561. Polytechnic Press, NY, 1962.
- [10] D. Caucal. Graphes canoniques de graphes algébriques. *ITA*, 24:339–352, 1990.
- [11] S. Christensen, H. Hüttel, and C. Stirling. Bisimulation equivalence is decidable for all context-free processes. *Information and Computation*, 121(2):143–148, 1995.
- [12] J.-C. Fernandez, L. Mounier, C. Jard, and T. Iron. On-the-fly verification of finite transition systems. *Formal Methods in System Design*, 1(2/3):251–273, 1992.
- [13] M. R. Henzinger, T. A. Henzinger, and P. W. Kopke. Computing simulations on finite and infinite graphs. In *Proc. FOCS*, pages 453–462. IEEE Computer Society, 1995.
- [14] Y. Hirshfeld, M. Jerrum, and F. Moller. A polynomial algorithm for deciding bisimilarity of normed context-free processes. *Theoretical Computer Science*, 158(1&2):143–159, 1996.
- [15] J. E. Hopcroft. An $n \log n$ algorithm for minimizing in a finite automaton. In *Proc. International Symposium of Theory of Machines and Computations*, pages 189–196. Academic Press, 1971.
- [16] J. E. Hopcroft and R. M. Karp. A linear algorithm for testing equivalence of finite automata. TR 114, Cornell Univ., December 1971.
- [17] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
- [18] H. Hosoya, J. Vouillon, and B. C. Pierce. Regular expression types for XML. *ACM Trans. Program. Lang. Syst.*, 27(1):46–90, 2005.
- [19] D. Lee and M. Yannakakis. Online minimization of transition systems (extended abstract). In *Proc. STOC*, pages 264–274. ACM, 1992.
- [20] O. Lengál, J. Simáček, and T. Vojnar. Vata: A library for efficient manipulation of non-deterministic tree automata. In *TACAS*, vol. 7214 of *LNCS*, pages 79–94. Springer, 2012.
- [21] M. Lenisa. From set-theoretic coinduction to coalgebraic coinduction: some results, some problems. *ENTCS*, 19:2–22, 1999.
- [22] D. Lucanu and G. Rosu. Circular coinduction with special contexts. In *Proc. ICFEM*, vol. 5885 of *LNCS*, pages 639–659. Springer, 2009.
- [23] A. Meyer and L. J. Stockmeyer. Word problems requiring exponential time. In *Proc. STOC*, pages 1–9. ACM, 1973.
- [24] R. Milner. *Communication and Concurrency*. Prentice Hall, 1989.
- [25] R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes, I/II. *Information and Computation*, 100(1):1–77, 1992.
- [26] D. Pous. Complete lattices and up-to techniques. In *Proc. APLAS*, vol. 4807 of *LNCS*, pages 351–366. Springer, 2007.
- [27] J. Rutten. Automata and coinduction (an exercise in coalgebra). In *Proc. CONCUR*, vol. 1466 of *LNCS*, pages 194–218. Springer, 1998.
- [28] D. Sangiorgi. On the bisimulation proof method. *Mathematical Structures in Computer Science*, 8:447–479, 1998.
- [29] D. Sangiorgi. *Introduction to Bisimulation and Coinduction*. Cambridge University Press, 2011.
- [30] A. Silva, F. Bonchi, M. Bonsangue, and J. Rutten. Generalizing the powerset construction, coalgebraically. In *Proc. FSTTCS*, vol. 8 of *LIPICs*, pages 272–283. Leibniz-Zentrum fuer Informatik, 2010.

- [31] D. Tabakov and M. Vardi. Experimental evaluation of classical automata constructions. In *Proc. LPAR*, vol. 3835 of *LNCS*, pages 396–411. Springer, 2005.
- [32] D. Turi and G. D. Plotkin. Towards a mathematical operational semantics. In *LICS*, pages 280–291, 1997.
- [33] M. D. Wulf, L. Doyen, T. A. Henzinger, and J.-F. Raskin. Antichains: A new algorithm for checking universality of finite automata. In *Proc. CAV*, vol. 4144 of *LNCS*, pages 17–30. Springer, 2006.

A. Smallest bisimulation and compositionality

In this appendix, we show some (unrelated) properties that have been discussed through the paper, but never formally stated.

The first property concerns the relation computed by $\text{Naive}(x, y)$. The following proposition shows that it is the *smallest bisimulation* relating x and y .

Proposition 8. *Let x and y be two states of a DFA. Let R_{Naive} be the relation built by $\text{Naive}(x, y)$. If $\text{Naive}(x, y) = \text{true}$, then R_{Naive} is the smallest bisimulation relating x and y , i.e., $R_{\text{Naive}} \subseteq R$, for all bisimulations R such that $(x, y) \in R$.*

Proof. We have already shown in Proposition 2 that R_{Naive} is a bisimulation. We need to prove that it is the smallest. Let R be a bisimulation such that $(x, y) \in R$. For all words $w \in A^*$ and pair of states (x', y') such that $x \xrightarrow{w} x'$ and $y \xrightarrow{w} y'$, it must hold that $(x', y') \in R$ (by definition of bisimulation).

By construction, for all $(x', y') \in R_{\text{Naive}}$ there exists a word $w \in A^*$, such that $x \xrightarrow{w} x'$ and $y \xrightarrow{w} y'$. Therefore all the pairs in R_{Naive} must be also in R , that is $R_{\text{Naive}} \subseteq R$. $\square$

The second property is

$$\llbracket X + Y \rrbracket = \llbracket X \rrbracket + \llbracket Y \rrbracket,$$

which we have used in the Introduction to give an intuition of bisimulation up to context and to show that the problem of language inclusion can be reduced to language equivalence. We believe that this property is interesting, since it follows from the categorical observation made in [30] that determinised NFA are bialgebras [32], like CCS processes. For this reason, we prove here that $\llbracket - \rrbracket : \mathcal{P}(S) \rightarrow 2^{A^*}$ is a semi-lattice homomorphism.

Theorem 7. *Let (S, o, t) be a non-deterministic automaton and $(\mathcal{P}(S), o^\sharp, t^\sharp)$ be the corresponding deterministic automaton obtained through the powerset construction. The function $\llbracket - \rrbracket : \mathcal{P}(S) \rightarrow 2^{A^*}$ is a semi-lattice homomorphism, that is, for all $X_1, X_2 \in \mathcal{P}(S)$,*

$$\llbracket X_1 + X_2 \rrbracket = \llbracket X_1 \rrbracket + \llbracket X_2 \rrbracket \quad \text{and} \quad \llbracket 0 \rrbracket = 0.$$

Proof. We prove that for all words $w \in A^*$, $\llbracket X_1 + X_2 \rrbracket(w) = \llbracket X_1 \rrbracket(w) + \llbracket X_2 \rrbracket(w)$, by induction on w .

- for ϵ , we have:

$$\begin{aligned} \llbracket X_1 + X_2 \rrbracket(\epsilon) &= o^\sharp(X_1 + X_2) \\ &= o^\sharp(X_1) + o^\sharp(X_2) = \llbracket X_1 \rrbracket(\epsilon) + \llbracket X_2 \rrbracket(\epsilon). \end{aligned}$$

- for $a \cdot w$, we have:

$$\begin{aligned} \llbracket X_1 + X_2 \rrbracket(a \cdot w) &= \llbracket t_a^\sharp(X_1 + X_2) \rrbracket(w) && \text{(by definition)} \\ &= \llbracket t_a^\sharp(X_1) + t_a^\sharp(X_2) \rrbracket(w) && \text{(by definition)} \\ &= \llbracket t_a^\sharp(X_1) \rrbracket(w) + \llbracket t_a^\sharp(X_2) \rrbracket(w) && \text{(by induction hypothesis)} \\ &= \llbracket X_1 \rrbracket(a \cdot w) + \llbracket X_2 \rrbracket(a \cdot w). && \text{(by definition)} \end{aligned}$$

For the second part, we prove that for all words $w \in A^*$, $\llbracket 0 \rrbracket(w) = 0$, again by induction on w . *Base case:* $\llbracket 0 \rrbracket(\epsilon) = o^\sharp(0) = 0$. *Inductive case:* $\llbracket 0 \rrbracket(a \cdot w) = \llbracket t_a^\sharp(0) \rrbracket(w) = \llbracket 0 \rrbracket(w)$ that by induction hypothesis is 0. $\square$

B. Proofs of Section 2

Proposition 1. Two states are language equivalent iff there exists a bisimulation that relates them.

Proof. Let $R_{\llbracket - \rrbracket}$ be the relation $\{(x, y) \mid \llbracket x \rrbracket = \llbracket y \rrbracket\}$. We prove that $R_{\llbracket - \rrbracket}$ is a bisimulation. If $x R_{\llbracket - \rrbracket} y$, then $o(x) = \llbracket x \rrbracket(\epsilon) = \llbracket y \rrbracket(\epsilon) = o(y)$. Moreover, for all $a \in A$ and $w \in A^*$, $\llbracket t_a(x) \rrbracket(w) = \llbracket x \rrbracket(a \cdot w) = \llbracket y \rrbracket(a \cdot w) = \llbracket t_a(y) \rrbracket(w)$ that means $\llbracket t_a(x) \rrbracket = \llbracket t_a(y) \rrbracket$, that is $t_a(x) R_{\llbracket - \rrbracket} t_a(y)$.

We now prove the other direction. Let R be a bisimulation. We want to prove that $x R y$ entails $\llbracket x \rrbracket = \llbracket y \rrbracket$, i.e., for all $w \in A^*$, $\llbracket x \rrbracket(w) = \llbracket y \rrbracket(w)$. We proceed by induction on w . For $w = \epsilon$, we have $\llbracket x \rrbracket(\epsilon) = o(x) = o(y) = \llbracket y \rrbracket(\epsilon)$. For $w = a \cdot w'$, since R is a bisimulation, we have $t_a(x) R t_a(y)$ and thus $\llbracket t_a(x) \rrbracket(w') = \llbracket t_a(y) \rrbracket(w')$ by induction. This allows us to conclude since $\llbracket x \rrbracket(a \cdot w') = \llbracket t_a(x) \rrbracket(w')$ and $\llbracket y \rrbracket(a \cdot w') = \llbracket t_a(y) \rrbracket(w')$. $\square$

Lemma 1. The following functions are compatible:

id : the identity function;

$f \circ g$: the composition of compatible functions f and g ;

$\bigcup F$: the pointwise union of an arbitrary family F of compatible functions: $\bigcup F(R) = \bigcup_{f \in F} f(R)$;

f^ω : the (omega) iteration of a compatible function f .

Proof. The first two points are straightforward;

For the third one, assume that F is a family of compatible functions. Suppose that $R \mapsto R'$; for all $f \in F$, we have $f(R) \mapsto f(R')$ so that $\bigcup_{f \in F} f(R) \mapsto \bigcup_{f \in F} f(R')$.

For the last one, assume that f is compatible; for all n , f^n is compatible because (a) $f^0 = id$ is compatible (by the first point) and (b) $f^{n+1} = f \circ f^n$ is compatible (by the second point and induction hypothesis). By definition $f^\omega = \bigcup_n f^n$ and thus, by the third point, f^ω is compatible. $\square$

Lemma 2. The following functions are compatible:

- the constant reflexive function: $r(R) = \{(x, x) \mid \forall x \in S\}$;
- the converse function: $s(R) = \{(y, x) \mid x R y\}$;
- the squaring function: $t(R) = \{(x, z) \mid \exists y, x R y R z\}$.

Proof. r : observe that the identity relation $Id = \{(x, x) \mid \forall x \in S\}$ is always a bisimulation, i.e., $Id \mapsto Id$. Thus for all R, R' $r(R) = Id \mapsto Id = r(R')$.

s : observe that the definition of progression is completely symmetric. Therefore, if $R \mapsto R'$, then $s(R) \mapsto s(R')$.

t : assume that $R \mapsto R'$. For each $(x, z) \in t(R)$, there exists y such that $(x, y) \in R$ and $(y, z) \in R$. By assumption, (1) $o'(x) = o'(y) = o'(z)$ and (2) for all $a \in A$, $t'_a(x) R' t'_a(y) R' t'_a(z)$, that is $t'_a(x) t(R') t'_a(z)$. $\square$

C. Proofs of Section 3

Lemma 4. For all relations R , the relation $\rightsquigarrow_R$ is convergent.

Proof. We have that $Z \rightsquigarrow_R Z'$ implies $|Z'| > |Z|$, where $|X|$ denotes the cardinality of the set X (note that $\rightsquigarrow_R$ is irreflexive). Since $|Z'|$ is bounded by $|S|$, the number of states of the NFA, the relation $\rightsquigarrow_R$ is strongly normalising. We can also check that whenever $Z \rightsquigarrow_R Z_1$ and $Z \rightsquigarrow_R Z_2$, either $Z_1 = Z_2$ or there is some Z' such that $Z_1 \rightsquigarrow_R Z'$ and $Z_2 \rightsquigarrow_R Z'$. Therefore, $\rightsquigarrow_R$ is convergent. $\square$

Lemma 13. The relation $\rightsquigarrow_R$ is contained in $c(R)$.

Proof. If $Z \rightsquigarrow_R Z'$ then there exists $(X, Y) \in (s \cup \text{id})(R)$ such that $Z = Z + X$ and $Z' = Z + Y$. Therefore $Z \xrightarrow{c(R)} Z'$ and, thus, $\rightsquigarrow_R$ is contained in $c(R)$. $\square$

Lemma 14. Let $X, Y \in \mathcal{P}(S)$, we have $(X + Y) \downarrow_R = (X \downarrow_R + Y \downarrow_R) \downarrow_R$.

Proof. Follows from confluence (Lemma 4) and from the fact that for all $Z, Z', U, Z \rightsquigarrow_R Z'$ entails $U + Z \rightsquigarrow_R U + Z'$. $\square$

Theorem 3. For all relations R , and for all $X, Y \in \mathcal{P}(S)$, we have $X \downarrow_R = Y \downarrow_R$ iff $(X, Y) \in c(R)$.

Proof. From right to left. We proceed by induction on the derivation of $(X, Y) \in c(R)$. The cases for rules r , s , and t are straightforward. For rule id , suppose that $X \xrightarrow{R} Y$, we have to show $X \downarrow_R = Y \downarrow_R$:

- if $X = Y$, we are done;
- if $X \subsetneq Y$, then $X \rightsquigarrow_R X + Y = Y$;
- if $Y \subsetneq X$, then $Y \rightsquigarrow_R X + Y = X$;
- if neither $Y \subseteq X$ nor $X \subseteq Y$, then $X, Y \rightsquigarrow_R X + Y$.

(In the last three cases, we conclude by confluence—Lemma 4.)

For rule u , suppose by induction that $X_i \downarrow_R = Y_i \downarrow_R$ for $i \in 1, 2$; we have to show that $(X_1 + Y_1) \downarrow_R = (X_2 + Y_2) \downarrow_R$. This follows by Lemma 14.

From left to right. By Lemma 13, we have $X \xrightarrow{c(R)} X \downarrow_R$ for any set X , so that $X \xrightarrow{c(R)} X \downarrow_R = Y \downarrow_R \xrightarrow{c(R)} Y$. $\square$

Lemma 5. The three algorithms require at most $1 + v \cdot |R|$ iterations, where $|R|$ is the size of the produced relation; moreover, this bound is reached whenever they return true.

Proof. At each iteration, one pair is extracted from *todo*. The latter contains one pair before entering the loop and v pairs are added to it every time that a pair is added to R . $\square$

Lemma 15. Let x and y be two states of a DFA. Let R_{Naive} and R_{HK} be relations computed by $\text{Naive}(x, y)$ and $\text{HK}(x, y)$, respectively. If $\text{Naive}(x, y) = \text{HK}(x, y) = \text{true}$, then $e(R_{\text{Naive}}) = e(R_{\text{HK}})$.

Proof. By the proof of Proposition 3, $e^\omega(R_{\text{HK}})$ is a bisimulation. Since e is idempotent, we have $e^\omega = e$ and thus $e(R_{\text{HK}})$ is a bisimulation; we can thus deduce by Proposition 8 that $R_{\text{Naive}} \subseteq e(R_{\text{HK}})$. Moreover, by definition of the algorithms, we have $R_{\text{HK}} \subseteq R_{\text{Naive}}$. Summarising,

$$R_{\text{HK}} \subseteq R_{\text{Naive}} \subseteq e(R_{\text{HK}})$$

It follows that $e(R_{\text{HK}}) = e(R_{\text{Naive}})$, e being monotonic and idempotent. $\square$

Lemma 6. Let R_{Naive} , R_{HK} , and R_{HKC} denote the relations produced by the three algorithms. We have

$$|R_{\text{HKC}}|, |R_{\text{HK}}| \leq m \quad |R_{\text{Naive}}| \leq m^2, \quad (2)$$

where $m \leq 2^n$ is the number of accessible states in the determinised NFA and n is the number of states of the NFA. If the algorithms returned true, we moreover have

$$|R_{\text{HKC}}| \leq |R_{\text{HK}}| \leq |R_{\text{Naive}}|. \quad (3)$$

Proof. For the first point, let PS denote the set of (determinised NFA) states accessible from the two starting states, so that $m = |PS| \leq 2^n$. Since $R_{\text{Naive}} \subseteq PS \times PS$, we deduce $|R_{\text{Naive}}| \leq m^2$. Since each pair added to R_{HK} merges two distinct equivalence classes in $e(R_{\text{HK}})$, we necessarily have $|R_{\text{HK}}| \leq m$ (the largest partition of PS has exactly m singletons). Similarly, each pair added to R_{HKC} merges at least two distinct equivalence classes in $c(R_{\text{HK}})$, so that we also have $|R_{\text{HKC}}| \leq m$.

For the second point, $|R_{\text{HK}}| \leq |R_{\text{Naive}}|$ follows from the fact that $R_{\text{HK}} \subseteq R_{\text{Naive}}$, by definition of the algorithms. The other inequality is less obvious.

By construction, $R_{\text{HKC}} \subseteq R_{\text{Naive}}$ and, since e is monotonic, $e(R_{\text{HKC}}) \subseteq e(R_{\text{Naive}}) = e(R_{\text{HK}})$ (the latter equality is given by Proposition 15). In particular, there are more equivalence classes in $e(R_{\text{HKC}})$ than in $e(R_{\text{HK}})$; using the same argument as above, we deduce that $|R_{\text{HKC}}| \leq |R_{\text{HK}}|$. $\square$

Lemma 8. Let X, Y be two sets of states; let $R_{\subseteq}$ and $R_{\supseteq}$ be the relations computed by $\text{HKC}(X+Y, Y)$ and $\text{HKC}(X+Y, X)$, respectively. If $R_{\subseteq}$ and $R_{\supseteq}$ are bisimulations up to congruence, then the following relation is a bisimulation up to congruence:

$$R_{=} = \{(X', Y') \mid (X' + Y', Y') \in R_{\subseteq} \text{ or } (X' + Y', X') \in R_{\supseteq}\}.$$

Proof. Let $(X', Y') \in R_{=}$ and suppose that $(X' + Y', Y') \in R_{\subseteq}$ (the other case is symmetric).

First notice that all pairs in $R_{\supseteq}$ necessarily have the shape $(t_w^\#(X+Y), t_w^\#(X))$, for some word w . Since $R_{\supseteq}$ is a bisimulation up to congruence, $c(R_{\supseteq})$ is a bisimulation. Since $(X+Y, X) \in c(R_{\supseteq})$ then, for all words w , $(t_w^\#(X+Y), t_w^\#(X)) \in c(R_{\supseteq})$ and thus $(X' + Y', X') \in c(R_{\supseteq})$ (we have $X' = t_w^\#(X)$ and $Y' = t_w^\#(Y)$ for some word w).

Since $c(R_{\subseteq})$ and $c(R_{\supseteq})$ are bisimulations containing $(X' + Y', Y')$ and $(X' + Y', X')$, it holds that:

1. $o^\#(X') = o^\#(X' + Y') = o^\#(Y')$;
2. for all a , $t_a^\#(X' + Y') \xrightarrow{c(R_{\supseteq})} t_a^\#(X')$ and $t_a^\#(X' + Y') \xrightarrow{c(R_{\subseteq})} t_a^\#(Y')$.

By Lemma 7, $t_a^\#(Y') \subseteq t_a^\#(X') \downarrow_{R_{\supseteq}}$ and $X' \subseteq t_a^\#(Y') \downarrow_{R_{\subseteq}}$ and since all the rewriting rules for $R_{\subseteq}$ and $R_{\supseteq}$ are also rewriting rules for $R_{=}$, then $t_a^\#(Y') \subseteq t_a^\#(X') \downarrow_{R_{=}}$ and $t_a^\#(X') \subseteq t_a^\#(Y') \downarrow_{R_{=}}$. By the first observation in the proof of Lemma 7, this means that $t_a^\#(X') \xrightarrow{c(R_{=})} t_a^\#(Y')$. $\square$

D. Proofs of Section 4

Proposition 4. For all sets X, Y , we have $\llbracket X \rrbracket \subseteq \llbracket Y \rrbracket$ iff there exists a simulation T such that for all $x \in X$, $x \xrightarrow{T} Y$.

Proof. Let $T_{[-]}$ be the relation $\{(x, Y) \mid \llbracket x \rrbracket \subseteq \llbracket Y \rrbracket\}$. We prove that $T_{[-]}$ is a simulation. If $x \ T_{[-]} \ Y$, then $o(x) = \llbracket x \rrbracket(\epsilon) \leq \llbracket Y \rrbracket(\epsilon) = o^\sharp(Y)$. Moreover, for all $a \in A$ $x' \in t_a(x)$ and $w \in A^*$, $\llbracket x' \rrbracket(w) \subseteq \llbracket x \rrbracket(a \cdot w) \subseteq \llbracket Y \rrbracket(a \cdot w) = \llbracket t_a^\sharp(Y) \rrbracket(w)$ that means $\llbracket x' \rrbracket \subseteq \llbracket t_a^\sharp(Y) \rrbracket$, that is $t_a(x) \ T_{[-]} \ t_a^\sharp(Y)$.

We now prove the other direction. Let T be a simulation. We want to prove that $x \ T \ Y$ entails $\llbracket x \rrbracket \subseteq \llbracket Y \rrbracket$, i.e., for all $w \in A^*$, $\llbracket x \rrbracket(w) \subseteq \llbracket Y \rrbracket(w)$. We proceed by induction on w . For $w = \epsilon$, we have $\llbracket x \rrbracket(\epsilon) = o(x) \leq o^\sharp(Y) = \llbracket Y \rrbracket(\epsilon)$. For $w = a \cdot w'$, since T is a simulation, we have $t_a(x) \ T \ t_a^\sharp(Y)$ and thus $\llbracket t_a(x) \rrbracket(w') \subseteq \llbracket t_a^\sharp(Y) \rrbracket(w')$ by induction. This allows us to conclude since $\llbracket x \rrbracket(a \cdot w') = \llbracket t_a(x) \rrbracket(w')$ and $\llbracket y \rrbracket(a \cdot w') = \llbracket t_a^\sharp(y) \rrbracket(w')$. $\square$

Definition 10. A function $f : \mathcal{P}(S \times \mathcal{P}(S)) \rightarrow \mathcal{P}(S \times \mathcal{P}(S))$ is *s-compatible* if it is monotone and for all relations $T, T' \subseteq S \times \mathcal{P}(S)$, $T \rightsquigarrow_s T'$ entails $f(T) \rightsquigarrow_s f(T')$.

Lemma 16. Any simulation T up to an s-compatible function f ($T \rightsquigarrow_s f(T)$) is contained in a simulation, namely $f^\omega(T)$.

Proof. Same proof as for Proposition 3. $\square$

Lemma 17. The upward closure function $\uparrow$ is s-compatible.

Proof. We assume that $T \rightsquigarrow_s T'$ and we prove that $\uparrow T \rightsquigarrow_s \uparrow T'$. If $x \uparrow T \ Y$, then $\exists Y' \subseteq Y$ such that $x \ T \ Y'$. Since $Y' \subseteq Y$, $o^\sharp(Y') \leq o^\sharp(Y)$ and $t_a^\sharp(Y') \subseteq t_a^\sharp(Y)$ for all $a \in A$. Since $T \rightsquigarrow_s T'$ and $x \ T \ Y'$, then $o(x) \leq o^\sharp(Y') \leq o^\sharp(Y)$ and $t_a(x) \uparrow T' \ t_a^\sharp(Y)$ for all $a \in A$. $\square$

Theorem 4. Any simulation up to $\uparrow$ is contained in a simulation.

Proof. By Lemmas 16 and 17. $\square$

Lemma 18. The relation

$$\begin{aligned} R' = & \{(x + y, z)\} \\ & + \{(x + Y_i + y_{i+1}, Z_{i+1}) \mid i < n\} \\ & + \{(x + Y_i + x_{i+1}, Z_{i+1}) \mid i < n\} \end{aligned}$$

is a bisimulation up to congruence for the NFA in Fig. 5.

Proof. First notice that

$$X_1 + y \ c(R') \ x + Y_1 \ c(R') \ Z_1$$

We then consider each kind of pair of R' separately:

- (x, y) : we have $o^\sharp(x + y) = 0 = o^\sharp(z)$ and $t_a^\sharp(x + y) = X_1 + y \ R' \ Z_1 = t_a^\sharp(z)$ and, similarly, $t_b^\sharp(x + y) = x + Y_1 \ R' \ Z_1 = t_b^\sharp(z)$.
- $(x + Y_i + y_{i+1}, Z_{i+1})$: both members are accepting iff $i + 1 = n$; setting $j = \min(i + 2, n)$, we have

$$\begin{aligned} t_a^\sharp(x + Y_i + y_{i+1}) &= X_1 + y + y_2 + \dots + y_j \\ c(R') \ x + Y_1 + y_2 + \dots + y_j \\ &= x + Y_j \ R' \ Z_j = t_a^\sharp(Z_{i+1}) \end{aligned}$$

and

$$t_b^\sharp(x + Y_i + y_{i+1}) = x + Y_j \ R' \ Z_j = t_b^\sharp(Z_{i+1})$$

- $(x + Y_i + x_{i+1}, Z_{i+1})$: both members are accepting iff $i + 1 = n$; if $i + 1 < n$ then we have:

$$\begin{aligned} t_a^\sharp(x + Y_i + x_{i+1}) &= X_1 + y + y_2 + \dots + y_{i+1} + x_{i+2} \\ c(R') \ x + Y_1 + y_2 + \dots + y_{i+1} + x_{i+2} \\ &= x + Y_{i+1} + x_{i+2} \\ R' \ Z_{i+2} &= t_a^\sharp(Z_{i+1}) \end{aligned}$$

and

$$t_b^\sharp(x + Y_i + x_{i+1}) = x + Y_{i+1} + x_{i+2} \ R' \ Z_{i+2} = t_b^\sharp(Z_{i+1})$$

otherwise, i.e., $i + 1 = n$, notice that:

$$\begin{aligned} x + Y_n + x_n \ c(R') \ Z_n + y_n \\ c(R') \ x + Y_n + y_n &= x + Y_n \\ c(R') \ Z_n &= t_a^\sharp(Z_n), \end{aligned}$$

from which we deduce:

$$\begin{aligned} t_a^\sharp(x + Y_i + x_n) &= X_1 + y + y_2 + \dots + y_n + x_n \\ c(R') \ x + Y_1 + y_2 + \dots + y_n + x_n \\ &= x + Y_n + x_n \ c(R') \ t_a^\sharp(Z_n) \end{aligned}$$

and

$$t_b^\sharp(x + Y_i + x_n) = x + Y_n + x_n \ c(R') \ t_b^\sharp(Z_n)$$

$\square$

E. Proofs of Section 5

Theorem 5. Any simulation up to $\hat{\lambda}$ is contained in a simulation.

Proof. By Lemma 16, it suffices to show that $\hat{\lambda}$ is s-compatible. Suppose that $T \rightsquigarrow_s T'$, we have to show that $\hat{\lambda}T \rightsquigarrow_s \hat{\lambda}T'$. Assume that $x \ \hat{\lambda}T \ Y$.

- if $x \preceq^{\forall\exists} Y$ then $x \preceq y$ for some $y \in Y$. Therefore, we have $o(x) \leq o(y) \leq o^\sharp(Y)$ and for all $a \in A$, $x' \in t_a(x)$, we have some $y' \in t_a(y)$ with $x' \preceq y'$. Since $t_a(y) \subseteq t_a^\sharp(Y)$, we deduce $x' \preceq^{\forall\exists} t_a^\sharp(Y)$, and hence $x' \ \hat{\lambda}T' \ t_a^\sharp(Y)$, as required.
- otherwise, we have some $(x', Y') \in T$ such that $(x', Y') \sqsubseteq^\preceq (x, Y)$, i.e., $x \preceq x'$ and $Y' \preceq^{\forall\exists} Y$. Since $T \rightsquigarrow_s T'$, we have $o(x) \leq o(x') \leq o^\sharp(Y') \leq o^\sharp(Y)$. Now take some $x'' \in t_a(x)$, we have some $x''' \in t_a(x')$ with $x'' \preceq x'''$, and since $T \rightsquigarrow_s T'$, we know $x''' \ T' \ t_a^\sharp(Y')$. It suffices to show that $t_a^\sharp(Y') \preceq^{\forall\exists} t_a^\sharp(Y)$ to conclude; this follows easily from $Y' \preceq^{\forall\exists} Y$ and from the definition of similarity. $\square$

Lemma 11. For all relations R , $c'(R) = c(R \cup \overline{\preceq})$.

Proof. The inclusion $c(R \cup \overline{\preceq}) \subseteq c'(R)$ is trivial. For the other inclusion we take $d = r' \cup s \cup t \cup u \cup id$ and we prove by induction that for all natural numbers n , $d^n(R) \subseteq c(R \cup \overline{\preceq})$. For $n = 0$, $d^0(R) = R \subseteq c(R \cup \overline{\preceq})$. For $n + 1$, $d^{n+1}(R) = d(d^n(R))$. By induction hypothesis, $d^n(R) \subseteq c(R \cup \overline{\preceq})$ and, by monotonicity of d , $d(d^n(R)) \subseteq d(c(R \cup \overline{\preceq}))$. By definition of d , the latter is equal to $c(R \cup \overline{\preceq})$. $\square$

Lemma 12. For all relations $T \subseteq S \times \mathcal{P}(S)$, $\widehat{\lambda}T \subseteq c'(\widehat{T})$.

Proof. If $(x + Y, Y) \in \widehat{\mathcal{L}}\widehat{T}$, then either (a) $x \preceq^{\forall\exists} Y$ or (b) there exists $x \preceq x'$ and $Y' \preceq^{\forall\exists} Y$ such that $(x', Y') \in T$. We have to show $(x+Y, Y) \in c'(\widehat{T})$, i.e., $(x+Y, Y) \in c(\widehat{T} + \preceq)$ by Lemma 11, that is $x \in Y \downarrow_{\widehat{T} + \preceq}$ by Lemma 7. For (b), we have:

$$\begin{aligned} Y &\rightsquigarrow_{\widehat{T} + \preceq}^* Y + Y' & (Y' \preceq^{\forall\exists} Y) \\ &\rightsquigarrow_{\widehat{T} + \preceq} Y + Y' + x' & ((x' + Y', Y') \in \widehat{T}) \\ &\rightsquigarrow_{\widehat{T} + \preceq} Y + Y' + x' + x & (x \preceq x') \end{aligned}$$

$x \in Y \downarrow_{\widehat{T} + \preceq}$ follows by confluence (Lemma 4). For (a), we immediately have that $Y \rightsquigarrow_{\widehat{T} + \preceq} Y + x$. $\square$

Proposition 7. If T is a simulation up to $\mathcal{L}$, then $\widehat{T}$ is a bisimulation up to c' .

Proof. First observe that if $T \rightsquigarrow_s T'$, then $\widehat{T} \rightsquigarrow u^\omega(\widehat{T}')$. Therefore, if $T \rightsquigarrow_s \uparrow T$, then $\widehat{T} \rightsquigarrow u^\omega(\uparrow \widehat{T})$. By Lemma 12, $\widehat{T} \rightsquigarrow u^\omega(c'(\widehat{T})) = c'(\widehat{T})$. $\square$

Lemma 19. *The relation*

$$\begin{aligned} R'' &= \{(z+x+y, x+y)\} \\ &\cup \{(Z_{i+1}+X_i+y+y_{i+1}, X_i+y+y_{i+1}) \mid i < n\} \\ &\cup \{(Z_{i+1}+X_{i+1}+y, X_{i+1}+y) \mid i < n\}, \end{aligned}$$

is a bisimulation up to c' for the NFA in Figure 6.

Proof. Let X'_i be the set X_i without x_1 and note that $X_i \xrightarrow{a} X_{i+1}$ and $X_i \xrightarrow{b} X'_{i+1}$. First we observe that for all i ,

$$X'_i + Y_1 \rightsquigarrow_{R'' \cup \preceq} X'_i + Y_1 + Z_1 \rightsquigarrow_{R'' \cup \preceq} X'_i + Y_1 + Z_1 + x_1$$

where the first reduction is given by $(Z_1 + X_0 + y + y_1, X_0 + y + y_1) \in R''$ and the second by $x_1 \preceq z_1$. Since $X'_i + x_1 = X_i$, then one can apply the third kind of pairs in R'' , so that

$$X'_i + Y_1 \rightsquigarrow_{R'' \cup \preceq}^* X_i + Y_1 + Z_i$$

that is $Z_i \subseteq (X'_i + Y_1) \downarrow_{R'' \cup \preceq}$. By Lemmas 7 and 11, this means that

$$Z_i + X'_i + Y_1 \xrightarrow{c'} (R'') X'_i + Y_1 \quad (2)$$

If we moreover have y_{i+1} , we can apply the second kind of pair in R'' and obtain

$$X'_i + Y_1 + y_{i+1} \rightsquigarrow_{R'' \cup \preceq}^* X_i + Y_1 + Z_{i+1} + y_{i+1}$$

that is

$$Z_{i+1} + X'_i + Y_1 + y_{i+1} \xrightarrow{c'} (R'') X'_i + Y_1 + y_{i+1} \quad (3)$$

With (2) and (3), it is easy to prove that R'' is a bisimulation up to c' , by simply proceeding by cases:

- $(z+x+y, x+y)$: we have $o^\sharp(x+y+z) = 0 = o^\sharp(x+y)$ and $t_a^\sharp(x+y+z) = Z_1 + X_1 + y \xrightarrow{R''} X_1 + y = t_a^\sharp(x+y)$ and, similarly, $t_b^\sharp(x+y+z) = Z_1 + x + Y_1 \xrightarrow{R''} x + Y_1 = t_b^\sharp(z)$.
- $(Z_{i+1}+X_i+y+y_{i+1}, X_i+y+y_{i+1})$ and $i < n - 1$: both members are not accepting;

$$\begin{aligned} t_a^\sharp(Z_{i+1}+X_i+y+y_{i+1}) &= Z_{i+2}+X_{i+1}+y+y_{i+2} \\ &\xrightarrow{R''} X_{i+1}+y+y_{i+2} \\ &= t_a^\sharp(X_i+y+y_{i+1}) \end{aligned}$$

and

$$\begin{aligned} t_b^\sharp(Z_{i+1}+X_i+y+y_{i+1}) &= Z_{i+2}+X'_{i+1}+Y_1+y_{i+2} \\ &\xrightarrow{c'(R'')} X'_{i+1}+Y_1+y_{i+2} \\ &= t_b^\sharp(X_i+y+y_{i+1}) \end{aligned}$$

- $(Z_n+X_{n-1}+y+y_n, X_{n-1}+y+y_n)$ and $i = n - 1$: both members are accepting;

$$\begin{aligned} t_a^\sharp(Z_n+X_{n-1}+y+y_n) &= Z_n+X_n+y \\ &\xrightarrow{R''} X_n+y \\ &= t_a^\sharp(X_{n-1}+y+y_n) \end{aligned}$$

and

$$\begin{aligned} t_b^\sharp(Z_n+X_{n-1}+y+y_n) &= Z_n+X'_n+Y_1 \\ &\xrightarrow{c'(R'')} X'_n+Y_1 \\ &= t_b^\sharp(X_{n-1}+y+y_n) \end{aligned}$$

- $(Z_{i+1}+X_{i+1}+y, X_{i+1}+y)$ and $i < n - 1$: both members are not accepting;

$$\begin{aligned} t_a^\sharp(Z_{i+1}+X_{i+1}+y) &= Z_{i+2}+X_{i+2}+y \\ &\xrightarrow{R''} X_{i+2}+y \\ &= t_a^\sharp(X_{i+1}+y) \end{aligned}$$

and

$$\begin{aligned} t_b^\sharp(Z_{i+1}+X_{i+1}+y) &= Z_{i+2}+X'_{i+2}+Y_1 \\ &\xrightarrow{c(R'')} X'_{i+2}+Y_1 \\ &= t_b^\sharp(X_{i+1}+y) \end{aligned}$$

- (Z_n+X_n+y, X_n+y) : both members are accepting; Moreover,

$$\begin{aligned} t_a^\sharp(Z_n+X_n+y) &= Z_n+X_n+y \\ &\xrightarrow{R''} X_n+y = t_a^\sharp(X_n+y) \end{aligned}$$

and

$$\begin{aligned} t_b^\sharp(Z_n+X_n+y) &= Z_n+X'_n+Y_1 \\ &\xrightarrow{c(R'')} X'_n+Y_1 \\ &= t_b^\sharp(X_n+y) \end{aligned}$$

The cases for the letter c are always trivial since $Z_i \xrightarrow{c} 0$. $\square$