A LOWER BOUND ON WEB SERVICES COMPOSITION

ANCA MUSCHOLL[a] AND IGOR WALUKIEWICZ[b]

[a] LaBRI, Université Bordeaux, 351, Cours de la Libération, F-33 405, Talence cedex, France
e-mail address: anca@labri.fr

[b] CNRS LaBRI, 351, Cours de la Libération, F-33 405, Talence cedex, France
e-mail address: igw@labri.fr

ABSTRACT. A web service is modeled here as a finite state machine. A composition problem for web services is to decide if a given web service can be constructed from a given set of web services; where the construction is understood as a simulation of the specification by a fully asynchronous product of the given services. We show an EXPTIME-lower bound for this problem, thus matching the known upper bound. Our result also applies to richer models of web services, such as the Roman model.

1. INTRODUCTION

Inherently distributed applications such as web services [1] increasingly get into the focus of automated verification techniques. Often, some basic e-services are already implemented, but no such simple service can answer to a more complex query. For instance, a user interested in hiking Mt. Everest will ask a travel agency for information concerning weather forecast, group travels, guides etc. The travel agency will contact different e-services, asking for such information and making appropriate reservations, if places are available. In general, single services such as weather forecast or group reservations, are already available and it is important to be able to reuse them without any change. The task of the travel agency is to compose basic e-services in such a way that the user's requirements are met (and eventually some constraints wrt. the called services, such as avoiding unreliable ones). Thus, one main objective is to be able to check automatically that the composition of basic e-services satisfies certain desirable properties or realizes another complex e-service.

In this paper we study a problem that arises in the *composition* of e-services as considered in [2, 3, 4]. The setting is the following: we get as input a specification (goal) $\mathcal{B}$, together with n available services $\mathcal{A}_1, \ldots, \mathcal{A}_n$. Then we ask whether the composition of the services $\mathcal{A}_i$ can simulate the behavior of the goal $\mathcal{B}$. This problem is known as *composition synthesis*. It amounts to synthesize a so-called *delegator*, that tells at any moment which

service must perform an action. In essence, a delegator implements a simulation relation of the goal service $\mathcal{B}$ by the composition of the available services $\mathcal{A}_i$. In the most general setting, as considered for instance in [9, 8, 7], services are modeled by communicating state machines [5], that have access to some local data. In this paper, we reconsider the simplified setting of the so-called Roman model [2] where services are finite state processes with no access to data and no mutual synchronization. This restriction is severe, however sufficient for our purposes, since our primary motivation is to obtain a complexity lower bound for the composition synthesis problem.

In this paper we study the complexity of the composition synthesis problem in the very simple setting where the composition of the finite state machines $\mathcal{A}_i$ is fully asynchronous (in particular there is no communication). This case is interesting for two reasons. It is known to be decidable in EXPTIME [2], contrary to some richer frameworks where it is undecidable [3]. It is also probably the simplest setting where the problem can be formulated, thus the complexity of this variant gives a lower bound on the complexity of any other variants of the synthesis problem. A related problem arises when instead of simulation one considers bisimulation. This is sometimes called *orchestration problem*, where the issue is to find a communication architecture of the available services, that is equivalent to the goal, modulo bisimulation. In our setting, this problem amounts to checking if the asynchronous composition of finite state machines is bisimilar to a given machine.

The main result of this paper is the EXPTIME lower bound for the composition synthesis problem. We also show that the same question can be solved in polynomial time if we assume that the sets of actions of the available machines are pairwise disjoint, i.e., each request can be handled by precisely one service. Note that in the latter case, the set of actions depends on the number of processes, whereas for the first result we show that the case where the set of actions is fixed is already EXPTIME-hard. We also show that the orchestration (bisimulation) problem is NLOGSPACE complete, independently of whether the sets of actions of the components are disjoint or not[1]. This result, however, is less interesting the context of service composition. The bisimulation requirement means that that the client (goal automaton) should be prepared to admit all possible interleavings in the composition, which usually makes the specification too complex.

Similar kinds of questions were also considered by the verification community. There is a large body of literature on the complexity of bisimulation and simulation problems for different kinds of process calculi (for a survey see [12]). A result that is most closely related to ours is the EXPTIME completeness of simulation and bisimulation between non-flat systems [10]. The main difference to our setting is that there both a system and services are given as composition of finite state machines using (binary) *synchronization* on actions, i.e., an action can synchronize two services. In a sense this paper shows that the lower bound for the simulation holds even without any synchronization.

This paper is an extended version of the conference publication [11]. In particular, the characterization of the complexity of the bisimulation problem is new.

---

[1]This problem is easier than checking bisimulation between a BPP and a finite state automaton, which is P-complete. The reason is that the finite-state automaton is deterministic in our setting.

2. Notations

We denote throughout this paper tuples of states (i.e., global states of a product automaton) by bold characters $\vec{q}, \vec{s}, \vec{t}, \ldots$. Unless otherwise stated, the components of vector $\vec{t}$ are $t_1, \ldots, t_n$.

An *asynchronous* product of n deterministic automata

$$\mathcal{A}_i = \langle Q_i, \Sigma_i, q_i^0, \delta_i : Q_i \times \Sigma_i \to Q_i \rangle$$

is a nondeterministic automaton:

$$\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n = \langle Q, \Sigma, \vec{q}, \delta : Q \times \Sigma \to \mathcal{P}(Q) \rangle$$

where: $Q = Q_1 \times \cdots \times Q_n$; $\Sigma = \bigcup_{i=1,\ldots,n} \Sigma_i$; $\vec{q} = (q_1^0, \ldots, q_n^0)$; and δ is defined by:

$\vec{t} \in \delta(\vec{s}, a)$ iff for some i, $t_i = \delta_i(s_i, a)$ and for all $j \neq i$ we have $t_j = s_j$.

Observe that the product automaton can be non deterministic because the alphabets Σ_i are not necessarily disjoint.

We define a *simulation relation* on nondeterministic automata in a standard way. Take two nondeterministic automata $\mathcal{A} = \langle Q_A, \Sigma, q_A^0, \delta_A : Q_A \times \Sigma \to \mathcal{P}(Q_A) \rangle$ and $\mathcal{B} = \langle Q_B, \Sigma, q_B^0, \delta_B : Q_B \times \Sigma \to \mathcal{P}(Q_B) \rangle$ over the same alphabet. The simulation relation $\preccurlyeq \subseteq Q_A \times Q_B$ is the biggest relation such that if $q_A \preccurlyeq q_B$ then for every $a \in \Sigma$ and every $q_A' \in \delta_A(q_A, a)$ there is $q_B' \in \delta_B(q_B, a)$ such that $q_A' \preccurlyeq q_B'$. We write $\mathcal{A} \preccurlyeq \mathcal{B}$ if $q_A^0 \preccurlyeq q_B^0$.

Problem: Given n deterministic automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ and a deterministic automaton $\mathcal{B}$ decide if $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.

We will show that this problem is EXPTIME-complete. It is clearly in EXPTIME as one can construct the product $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ explicitly and calculate the biggest simulation relation with $\mathcal{B}$. The rest of this paper will contain the proof of EXPTIME-hardness. We will start with the PSPACE-hardness, as this will allow us to introduce the method and some notation.

3. A Pspace lower bound

We will show PSPACE-hardness of the problem by reducing it to the existence of a looping computation of a linearly space bounded deterministic Turing machine. The presented proof of the PSPACE bound has the advantage to generalize to the encoding of alternating machines that we will present in the following section.

Fix a deterministic Turing machine M working in space bounded by the size of its input. We want to decide if on a given input the computation of the machine loops. Thus we do not need any accepting states in the machine and we can assume that there are no transitions from rejecting states. We denote by Q the states of M and by Γ the tape alphabet of M. A *configuration* of M is a word over $\Gamma \cup (Q \times \Gamma)$ with exactly one occurrence of a letter from $Q \times \Gamma$. A configuration is of size n if it is a word of length n. Transitions of M will be denoted as $qa \longrightarrow q'bd$, where q, q' are the old/new state, a, b the old/new tape symbol and $d \in \{l, r\}$ the left/right head move (w.l.o.g. we assume that M moves the head in each step).

Suppose that the input is a word w of size n. We will construct automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ and $\mathcal{B}$ such that $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ iff the computation of M on w is infinite.

We start with some auxiliary alphabets. For every $i = 1, \ldots, n$ let
$$\Gamma_i = \Gamma \times \{i\} \quad \text{and} \quad \Delta_i = (Q \times \Gamma_i) \cup (Q \times \Gamma_i \times \{l, r\}).$$
We will write a_i instead of (a, i) for elements of Γ_i. Let also $\Delta = \bigcup_{i=1,\ldots,n} \Delta_i$.

The automaton $\mathcal{A}_i = \langle Q_i, \Sigma_i, q_i^0, \longrightarrow \rangle$ is defined as follows:

- The set of states is $Q_i = \Gamma \cup (Q \times \Gamma) \cup \{\top\}$, and the alphabet of the automaton is $\Sigma_i = \Delta$.
- We have transitions:
  - $a \xrightarrow{qa_i} qa$, for all $a \in \Gamma$ and $q \in Q$,
  - $qa \xrightarrow{q'b_i d} b$, for $qa \to q'bd$ the transition of M on qa (there is at most one).
  - From a, transitions on letters in $\Delta_i \setminus \{qa_i : q \in Q\}$ go to $\top$. Similarly, from qa transitions on $\Delta_i \setminus \{qb_i d\}$ go to $\top$ if there is a transition of M on qa; if not, then qa has no outgoing transitions. From $\top$ there are self-loops on all letters from Δ.
- For $i = 2, \ldots, n$ the initial state of $\mathcal{A}_i$ is w_i, the i-th letter of w; for $\mathcal{A}_1$ the initial state is $q^0 w_1$, i.e., the initial state of M and the first letter of w.

Figure 1 shows a part of $\mathcal{A}_i$:

Figure 1: Part of $\mathcal{A}_i$

The idea is classical: automaton $\mathcal{A}_i$ controls the i-th tape symbol, whereas automaton $\mathcal{B}$ defined below is in charge of the control part of M. The challenge is to do this without using any synchronization between adjacent automata $\mathcal{A}_i, \mathcal{A}_{i+1}$. Next, we introduce an automaton K that will be used to define $\mathcal{B}$ (see also Figure 2). The set of states of K is $Q_K = \{s, e\} \cup (Q \times \bigcup \Gamma_i \times \{l, r\})$; the initial state is s and the final one e; the alphabet is Δ; the transitions are defined by:

- $s \xrightarrow{q'b_i r} q'b_i r$ for $i = 1, \ldots, n-1$, whenever we have a transition $qa \to q'br$ in M for some state q and some letter a;
- $s \xrightarrow{q'b_{i+1} l} q'b_{i+1} l$ for $i = 1, \ldots, n-1$, whenever we have a transition $qa \to qbl$ in M for some state q and some letter a;
- $q'b_i r \xrightarrow{q'c_{i+1}} e$ and $q'b_{i+1} l \xrightarrow{q'c_i} e$ for all $c \in \Gamma$.

Figure 2 presents a schema of the automaton K. We define $\mathcal{B}$ as the deterministic automaton recognizing $(L(K))^*$, that is obtained by gluing together the states s and e.

Remark 1. *All $\mathcal{A}_i$ and $\mathcal{B}$ are deterministic automata of size polynomial in n. The input alphabets of the $\mathcal{A}_i$ are almost pairwise disjoint: the only states with common labels on outgoing transitions are the $\top$ states.*

Figure 2: Automaton K

Definition 3.1. We say that a configuration C of size n of M *corresponds* to a global state $\vec{s}$ of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ iff $s_i = C(i)$ for $i = 1, \ldots, n$; in other words, if the state of $\mathcal{A}_i$ is the same as the i-th letter of C.

Definition 3.2. We say that a global state $\vec{s}$ of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ is *proper* when there is no $\top$-state in $\vec{s}$.

Lemma 3.3. *If $\vec{s}$ is a proper state, then for every letter $a \in \Delta$ the automaton $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ has in state $\vec{s}$ at most one outgoing a-transition. Once the automaton enters a state that is not proper, it stays in non proper states.*

It is easy to see that from a non proper state, $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ can simulate any state of $\mathcal{B}$. The reason is that from $\top$, any move on letters from Δ is possible.

Lemma 3.4. *Suppose that $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ is in a state $\vec{s}$ that corresponds to a configuration C of M.*
- *If C is a configuration with no successor, then there is a word $v \in L(K)$ that cannot be simulated by $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ from $\vec{s}$.*
- *Otherwise, the successor configuration $C \vdash C'$ exists, and there is a unique word $v \in L(K)$ such that $\vec{s} \xrightarrow{v} \vec{t}$ and $\vec{t}$ is proper. Moreover $\vec{t}$ corresponds to C'. All other words from $L(K)$ lead from $\vec{s}$ to non proper states of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.*

Proof. For the first claim, assume that $\vec{s}$ corresponds to a configuration, thus there is exactly one i such that $\mathcal{A}_i$ is in a state from $Q \times \Gamma$. The other automata are in states from Γ.

If C is terminal then $\mathcal{A}_i$ is in a state qa which has no outgoing transition. This means that this state can simulate no move on letters $q'b_i r$, for $q' \in Q$ and $b_i \in \Gamma_i$ (and such a move exists in K, as the machine M must have a move to the right if it is nontrivial). All other automata are also not capable to simulate $q'b_i r$ as they can do only moves on letters Δ_j for $j \neq i$.

Now suppose that $C \vdash C'$. To avoid special, but simple, cases suppose that the position i of the state is neither the first nor the last. Let $s_i = qa$ and suppose also that $qa \to q'br$ is the move of M on qa. The case when the move is to the left is similar.

The only possible move of K from s which will put $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ into a proper state is $q'b_i r$. This makes $\mathcal{A}_i$ to change the state to b and it makes K to change the state to $q'b_i r$. From this latter state the only possible move of K is on letters $q'c'_{i+1}$ for arbitrary $c' \in \Gamma$. Suppose that $\mathcal{A}_{i+1}$ is in the state $c = s_{i+1} \in \Gamma$, then all moves of K on $q'c'_{i+1}$ with $c' \neq c$ can be matched with a move to $\top$ of $\mathcal{A}_{i+1}$. On $q'c_{i+1}$ the automaton $\mathcal{A}_{i+1}$ goes to $q'c$ and automaton K goes to e. This way the state in the configuration is changed and transmitted to the right. We have that the new state of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ corresponds to the configuration C'.

□

Lemma 3.5. *We have $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ iff the computation of M on w is infinite.*

Proof. Recall that $\mathcal{B}$ is a deterministic automaton recognizing $(L(K))^*$, and has initial state s. The initial state of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ corresponds to the initial configuration C_0 of M on w. We show now for every state $\vec{t}$ corresponding to a configuration C of M: $s \preccurlyeq \vec{t}$ iff the computation of M starting in C is infinite.

From a configuration C, the machine M has only one computation: either infinite, or a finite one that is blocking. Suppose that the computation from C has at least one step and let C_1 be the successor configuration. By Lemma 3.4 from state s there is exactly one word $v_1 \in L(K)$ such that $\vec{t} \xrightarrow{v_1} \vec{t_1}$ in $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$, and $\vec{t_1}$ is proper. Moreover $\vec{t_1}$ corresponds to C_1. On all other words from $L(K)$, the product $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ reaches non proper states and from there it can simulate any future behaviour of $\mathcal{B}$. If C_1 has no successor configuration then, again by Lemma 3.4, there is a word in $L(K)$ that cannot be simulated by $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ from $\vec{t_1}$. If C_1 has a successor then we repeat the whole argument. Thus the behaviour of $\mathcal{B}$ from s can be simulated by $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ from the state corresponding to C iff the machine M has an infinite computation starting from C.

□

One can note that the construction presented in this section uses actions that are common to several processes in a quite limited way: the only states that have common outgoing labels are the $\top$ states from which all behaviours are possible. This observation motivates the question about the complexity of the problem when the automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ have pairwise disjoint alphabets. With this restriction, the simulation problem can be solved efficiently:

Theorem 3.6. *The following question can be solved in polynomial time:*
  Input: *n deterministic automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ over pairwise disjoint input alphabets, and a deterministic automaton $\mathcal{B}$.*
  Output: *decide if $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.*

Proof. Let $\mathcal{C}_i$ be a automaton with a single state $\top$, and with self-loops on every letter from the alphabet Σ_i of $\mathcal{A}_i$. We write $\mathcal{A}^{(i)}$ for the asynchronous product of all $\mathcal{C}_j$, $j \neq i$, and of $\mathcal{A}_i$. Similarly, $\vec{t}^{(i)}$ will denote $\vec{t}$ with all components but i replaced by $\top$. Suppose now that p is a state of $\mathcal{B}$, and $\vec{t}$ a state of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$. We write $p \preccurlyeq_i \vec{t}$ if p is simulated by $\vec{t}^{(i)}$ in $\mathcal{A}^{(i)}$. Notice that since $\mathcal{B}$ and $\mathcal{A}_i$ are both deterministic, we can decide if $p \not\preccurlyeq_i \vec{t}$ in logarithmic space (hence in polynomial time), by guessing simultaneously a path in $\mathcal{B}$ and one in $\mathcal{A}_i$.

We show now that $p \preccurlyeq \vec{t}$ in $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ iff $p \preccurlyeq_i \vec{t}$ for all i.

If $p \preccurlyeq \vec{t}$, then all the more $p \preccurlyeq \vec{t}^{(i)}$, since $\mathcal{C}_j$ can simulate $\mathcal{A}_j$ for all $j = 1, \ldots, n$. Conversely, assume that $p \preccurlyeq_i \vec{t}$ for all i, but $p \not\preccurlyeq \vec{t}$. This means that there exist computations $p \xrightarrow{a_1 \ldots a_k} p'$ in $\mathcal{B}$, $\vec{t} \xrightarrow{a_1 \ldots a_k} \vec{u}$ in $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ and a letter $a \in \Sigma_i$ for some i, such that p' has an outgoing a-transition, but $\vec{u}_i$ does not (in $\mathcal{A}_i$). Clearly, we also have a computation $\vec{t}^{(i)} \xrightarrow{a_1 \ldots a_k} \vec{u}^{(i)}$ in $\mathcal{A}^{(i)}$. Since $\vec{u}_i$ has no outgoing a-transition, so neither does $\vec{u}^{(i)}$, which contradicts $p \preccurlyeq_i \vec{t}$.

□

4. The complexity of simulation

This time we take an alternating Turing machine M working in space bounded by the size of the input. We want to decide if M has an infinite computation. This means that the machine can make choices of existential transitions in such a way that no matter what are the choices of universal transitions the machine can always continue. Clearly, one can reduce the word problem to this problem, hence it is Exptime-hard (see [6]; for more details on complexity see any standard textbook).

We will assume that M has always a choice between two transitions, i.e., for each non blocking state/symbol pair qa there will be precisely two distinct tuples $q'b'd'$, $q''b''d''$ such that $qa \to q'b'd'$ and $qa \to q''b''d''$. If q is existential then it is up to the machine to choose a move; if q is universal then the choice is made from outside. To simplify the presentation we will assume that $d' = d''$, i.e., both moves go in the same direction. Every machine can be transformed to an equivalent one with this property. We will also assume that the transitions are ordered in some way, so we will be able to say that $qa \to q'b'd$ is the first transition and $qa \to q''b''d$ is the second one.

Take the input word is w of size n. We will construct automata $\mathcal{A}_1', \mathcal{A}_1'', \ldots, \mathcal{A}_n', \mathcal{A}_n''$ and $\mathcal{B}$ such that $\mathcal{B}$ is simulated by $\mathcal{A}_1' \otimes \mathcal{A}_1'' \cdots \otimes \mathcal{A}_n' \otimes \mathcal{A}_n''$ iff there is an infinite alternating computation of M on w. The main idea is that automata $\mathcal{A}_i'$ and $\mathcal{A}_i''$ control the i-th tape symbol, as in the previous section, and each one is in charge of one of the two possible transitions (if any) when the input head is at position i in an existential state (universal moves are simpler).

We will modify a little the alphabets that we use. Let

$$\Delta_i' = (Q \times \Gamma_i) \cup (Q \times \Gamma_i \times \{l, r\} \times \{1\})$$
$$\Delta_i'' = (Q \times \Gamma_i) \cup (Q \times \Gamma_i \times \{l, r\} \times \{2\})$$

We then put $\Delta_i = \Delta_i' \cup \Delta_i''$, $\Delta = \bigcup_i \Delta_i$, $\Delta' = \bigcup_i \Delta_i'$ and $\Delta'' = \bigcup_i \Delta_i''$.

The automaton $\mathcal{A}_i'$ is defined as follows:

- The set of states is $Q_i' = \{\top\} \cup \Gamma \cup (Q \times \Gamma) \cup (Q \times \Gamma \times \{l, r\})$, the alphabet of the automaton is $\Sigma_i' = \Delta \cup \{\zeta\}$; where ζ is a new letter common to all automata.
- We have the following transitions:
  - $a \xrightarrow{qa_i} qa$ for all $a \in \Gamma$ and $q \in Q$,
  - $qa \xrightarrow{q'b_i'd1} b'$ and $qa \xrightarrow{q''b_i''d1} b''$ if q is an universal state and $qa \to q'b'd$, $qa \to q''b''d$ are the two transitions from qa. We have also transitions to $\top$ on all the letters from $\Delta_i' \setminus \{q'b_i'd1, q''b_i''d1\}$.
  - $qa \xrightarrow{\zeta} q'b'd \xrightarrow{q'b_i'd1} b'$ and $qa \xrightarrow{q''b_i''d1} b''$ if q is an existential state and $qa \to q'b'd$, $qa \to q''b''d$ are the first and the second transitions from qa, respectively. We have also transitions to $\top$ on all the letters from $\Delta_i' \setminus \{q''b_i''d1\}$. From $q'b'd$ all transitions on $\Delta_i' \setminus \{q'b_i'd1\}$ go to $\top$.
  - From a, transitions on letters in $\Delta_i' \setminus \{qa_i : q \in Q\}$ go to $\top$. If qa is terminal then there are no outgoing transitions from qa. From $\top$ there are self-loops on all letters from $\Delta^c := \Delta \cup \{\zeta\}$.
- The initial state of $\mathcal{A}_i'$ is w_i, the i-th letter of w except for $\mathcal{A}_1$ whose initial state is $q^0 w_1$, the initial state of M and the first letter of w.

Figure 3 below presents parts of $\mathcal{A}_i'$ corresponding to universal and existential states.
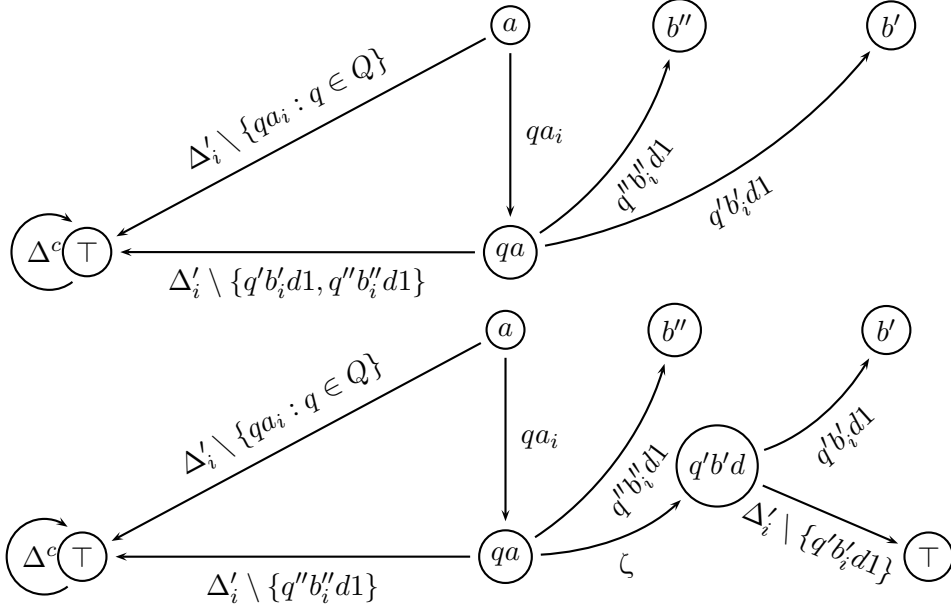
Figure 3: Parts of the automaton $\mathcal{A}'_i$ corresponding to universal and existential states q, respectively. The alphabet Δ^c is $\Delta \cup \{\zeta\}$.

The automaton $\mathcal{A}''_i$ is the same as $\mathcal{A}'_i$ with the difference that we replace every label $q''b''d1$ by $q'b'd2$, every $q'b'd1$ by $q''b''d2$ (notice the change of primes and double primes), every Δ'_i by Δ''_i and Δ' by Δ''. Moreover, state labels b' and b'' are exchanged, and state $q'b'd$ is relabeled $q''b''d$.

Next, we define a new automaton K that will be used to define new automaton $\mathcal{B}$. The states of K are

$$Q_K = \{s, e, choice\} \cup (Q \times \bigcup_i \Gamma_i \times \{l, r\})$$

plus some auxiliary states to implement transitions on two letters at a time. We will write transitions with two letters on them for readability. The initial state is s and the final one is e. The alphabet is $\Sigma_K = \bigcup \Sigma_i$. The transitions are defined by (cf. Figure 4):

- $s \xrightarrow{\zeta} choice$;
- $s \xrightarrow{(q'b_i r1)(q'b_i r2)} q'b_i r$ whenever we have a transition $qa \to q'br$ in M for some universal state q and some letter a, and similarly from *choice* instead of s when q is existential;
- $s \xrightarrow{(q'b_{i+1} l1)(q'b_{i+1} l2)} q'b_{i+1}l$ whenever we have a transition $qa \to q'bl$ in M for some universal state q and some letter a, and similarly from *choice* instead of s when q is existential;
- $q'b_i r \xrightarrow{(q'c_{i+1})^2} e$ and $q'b_{i+1}l \xrightarrow{(q'c_i)^2} e$ for all $c \in \Gamma$.

We define $\mathcal{B}$ as the deterministic automaton recognizing $(L(K))^*$ that is obtained by gluing together states s and e.

Remark 2. *All $\mathcal{A}'_i$, $\mathcal{A}''_i$ and $\mathcal{B}$ are deterministic and of size polynomial in n.*

Definition 4.1. A configuration C of size n *corresponds* to a global state $\vec{s}$ of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ if $s_{2i} = s_{2i-1} = C(i)$ for $i = 1, \ldots, n$; in other words, if the states of $\mathcal{A}'_i$ and $\mathcal{A}''_i$ are the same as the i-th letter of C.

Figure 4: Automaton K

Definition 4.2. We say that a global state $\vec{s}$ of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ is *proper* when $\top$ does not appear in $\vec{s}$.

It is easy to see that from a non proper state, $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ can simulate any state of $\mathcal{B}$. The reason is that from $\top$, any move on letters from Δ^c is possible.

Lemma 4.3. *Suppose that $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ is in a state $\vec{s}$ corresponding to a configuration C of M. If C has no successor configuration then there is a word $v \in L(K)$ that cannot be simulated by $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ from $\vec{s}$. Otherwise, C has two successor configurations $C \vdash C'$ and $C \vdash C''$. We have two cases:*

- *If C is universal then there are two words v' and v'' in $L(K)$: each leading from $\vec{s}$ to a unique state $\vec{t'}$ and $\vec{t''}$, respectively. These two states are proper and correspond to C' and C'', respectively. On all other words from $L(K)$, non proper states can be reached from $\vec{s}$.*
- *If C is existential, then on the letter ζ exactly two states are reachable from $\vec{s}$, call them $\vec{s'}$ and $\vec{s''}$. There is a word v' such that $\zeta v' \in L(K)$ and on v' from $\vec{s'}$ a unique state is reachable. This state is proper and corresponds to C'. Similarly there is a word v'' for $\vec{s''}$ and C''. On all words from $L(K)$ that are different from $\zeta v'$ and $\zeta v''$, non proper states can be reached from $\vec{s}$.*

Proof. As $\vec{s}$ corresponds to the configuration C, there is some i such that both automata $\mathcal{A}'_i$ and $\mathcal{A}''_i$ are in state qa, for some $q \in Q$ and $a \in \Gamma$, and all other automata are in states from Γ.

If C is a configuration without successor, then the state qa in $\mathcal{A}'_i$ and $\mathcal{A}''_i$ does not have any outgoing transition. Thus these automata cannot simulate the ζ transition of K from s. No other automaton $\mathcal{A}'_j$, or $\mathcal{A}''_j$ can simulate the ζ transition either, as they are all in states from Γ.

Suppose that C is an universal configuration with two possible transitions to the right, $qa \to q'b'r$ and $qa \to q''b''r$. The case when the moves are to the left is similar. In $\mathcal{A}'_i$ from the state qa we have a transition on $q'b'_i r1$ leading to b' and on $q''b''_i r1$ leading to b''. Similarly for $\mathcal{A}''_i$, but on $q'b'_i r2$ and $q''b''_i r2$. These transitions can simulate both transitions $(q'b'_i r1)(q'b'_i r2)$ and $(q''b''_i r1)(q''b''_i r2)$ that are possible from s in K. (All other transitions from s in K lead from $\vec{s}$ to a non proper state of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$.) Let us focus only

on the first case, when $(q'b'_ir1)(q'b'_ir2)$ is executed in K and the state $q'b'_ir$ is reached. From this state only transitions $(q'c'_{i+1})^2$ are possible, for all $c' \in \Gamma$. Suppose that $\mathcal{A}'_{i+1}$ and $\mathcal{A}''_{i+1}$ are in state $c \in \Gamma$. Transition $(q'c_{i+1})^2$ of K is simulated by moves to $q'c$ in both $\mathcal{A}'_{i+1}$ and $\mathcal{A}''_{i+1}$. This way the new state is transferred to the right. Transitions $(q'c'_{i+1})^2$ where $c \neq c'$ are simulated in $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ by moves of $\mathcal{A}'_{i+1}$ and $\mathcal{A}''_{i+1}$ to $\top$.

Suppose that C is an existential configuration, with possible transitions $qa \to q'b'r$ and $qa \to q''b''r$. The case when moves are to the left is similar. Consider first the transition of K from s that corresponds to the letter ζ. Both $\mathcal{A}'_i$ and $\mathcal{A}''_i$ can simulate this transition: the first goes to state $q'b'r$, and the second goes to $q''b''r$. Assume that it is the transition of $\mathcal{A}'_i$ that is taken; the other case is symmetric. We get to the position when K is in the state *choice*, $\mathcal{A}'_i$ is in the state $q'b'r$ and $\mathcal{A}''_i$ in the state qa. From *choice*, automaton K can do $(q'b'_ir1)(q'b'_ir2)$ that can be simulated by the transitions of $\mathcal{A}'_i$ and $\mathcal{A}''_i$ (every other transition of K can be simulated by a move of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ to a non proper state). Both automata reach the state b'. Automaton K is now in state $q'b_ir$ from where it can do $(q'c_{i+1})^2$ for any $c \in \Gamma$. The result of simulating these transitions while reaching a proper state is the transfer of the state to the right, in the same way as in the case of the universal move. Finally, it remains to see what happens if K makes a move from s that is different from ζ. In this case, at least one of the automata $\mathcal{A}'_i$, $\mathcal{A}''_i$ can simulate the corresponding transition on (pe_id1), (pe_id2) respectively, by going to state $\top$, since we suppose that in any configuration of M, the two outgoing transitions are distinct. Hence, a non proper state can be reached. $\square$

Theorem 4.4. *The following problem is* EXPTIME-*complete:*
   Input: *deterministic automata* $\mathcal{A}_1, \ldots, \mathcal{A}_n$ *and a deterministic automaton* $\mathcal{B}$.
   Output: *decide if* $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.

Proof. The problem is clearly in EXPTIME as the state space of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ can be constructed in EXPTIME. For EXPTIME hardness, we take an alternating machine M as at the beginning of this section and use the construction presented above together with Lemma 4.3. Recall, that $\mathcal{B}$ is a deterministic automaton obtained from the automaton K by gluing states s and e (cf. Figure 4). We also have that the initial state of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ corresponds to the initial configuration of M (in a way required by Definition 4.1). We will show that for every state $\vec{t}$ corresponding to a configuration C of M: $s \preccurlyeq \vec{t}$ iff M has an infinite alternating computation from C.

Consider a game of two players: Computer and Environment. Positions of the game are configurations of M. In existential configurations Computer chooses a successor configuration (with respect to the transition table of M). In universal configurations Environment makes a choice. Having an infinite alternating computation from C is equivalent to saying that in this game Computer has a strategy to avoid being blocked. At the same time, not having such a computation from C is equivalent to saying Environment has a strategy to reach a configuration with no successors. As this is a reachability game, for each such C there is a bound d_C (distance) on the number of steps in which Environment can force Computer into a blocking configuration. This distance is 0 if C is blocking; it is one plus the maximum over distances for two successor configurations if C is existential, and it is one plus the minimum over the distances of successor configurations if C is universal. (Here we assume that the distance is ∞ if Environment cannot win from C.).

Going back to the proof of the theorem, consider first the case when M does not have an infinite alternating computation from C. Let $\vec{t}$ be the state of $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$

corresponding to C. We show that $s \not\preccurlyeq \vec{t}$ by induction on the distance d_C. There are three possible cases:

- If $d_c = 0$ then is no transition possible from C. In this case Lemma 4.3 gives us an execution of $\mathcal{B}$ from s that cannot be simulated by $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$ from $\vec{t}$.
- If C is universal, there is a successor C_1 such that $d_C > d_{C_1}$. We take the word $v \in L(K)$ given by Lemma 4.3. The only way to simulate this word from $\vec{t}$ leads to the proper state $\vec{t}_1$ corresponding to C_1. By induction hypothesis $s \not\preccurlyeq \vec{t}_1$.
- If C is existential, then for both successor configurations, C' and C'', the distance is smaller. We make $\mathcal{B}$ execute ζ and then, depending how it was matched by $\mathcal{A}'_1 \otimes \mathcal{A}''_1 \cdots \otimes \mathcal{A}'_n \otimes \mathcal{A}''_n$, a word forcing the automaton to go to a proper state corresponding either to C' or to C''. Using the induction hypothesis we get that the simulation is not possible from s and the obtained states.

The case when M has an infinite alternating computation from C is very similar. In this case $d_C = \infty$. The means that if C is an existential computation then one of the successor configurations has distance equal to ∞. By Lemma 4.3 we can match ζ so that we go to the state corresponding to that configuration. If C is universal then both successor configurations have distance equal to ∞. Once again Lemma 4.3, tells us how to match every word from $L(K)$. $\square$

We conclude the section by showing that Theorem 4.4 still holds under the assumption that the alphabet of the automata $\mathcal{A}_i$ and $\mathcal{B}$ is of constant size.

Theorem 4.5. *Let Σ be a fixed alphabet of at least 2 letters. The following problem is* ExpTIME-*complete:*

Input: *deterministic automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ and a deterministic automaton $\mathcal{B}$ over the input alphabet Σ.*

Output: *decide if $\mathcal{B} \preccurlyeq \mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.*

Proof. We reduce directly from Theorem 4.4. Suppose that the input alphabet of all automata $\mathcal{A}_i, \mathcal{B}$ is $\Sigma \times \{1, \ldots, m\}$, for some m. Moreover, let S be the set of states of $\mathcal{B}$ and let $Q = Q_1 \times \cdots \times Q_n$ be the set of global states of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$.

In each automaton $\mathcal{A}_i, \mathcal{B}$ we replace every transition $s \xrightarrow{a_l} t$ by a sequence of transitions with labels from $\Sigma \cup \{\#, \$\}$ as follows:

$$s \xrightarrow{a} (stl0) \xrightarrow{\#} (stl1) \xrightarrow{\#} (stl2) \cdots \xrightarrow{\#} (stll) \xrightarrow{\$} t$$

The $(l+1)$ states $(stl0), \ldots, (stll)$ are new. Let $\mathcal{A}'_i, \mathcal{B}'$ be the automata obtained from $\mathcal{A}_i$, $\mathcal{B}$, with state space Q' and S', respectively.

Take $\preccurlyeq$, the largest simulation relation from $\mathcal{B}$ to $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$. We show how to extend $\preccurlyeq$ to $\preccurlyeq'$ such that $\preccurlyeq'$ is a simulation relation from $\mathcal{B}'$ to $\mathcal{A}'_1 \otimes \cdots \otimes \mathcal{A}'_n$ (not necessarily the largest one). Let $\preccurlyeq'$ be the union of $\preccurlyeq$ with the set of all pairs $((stlk), \vec{u}')$, where $s, t \in S$, $\vec{u}' = (u'_1, \ldots, u'_n) \in Q'$, and such that:

- $s \xrightarrow{a_l} t$ and $\vec{v} \xrightarrow{a_l} \vec{w}$ for some $a \in \Sigma$, $\vec{v} = (v_1, \ldots, v_n)$ and $\vec{w} = (w_1, \ldots, w_n)$ such that $s \preccurlyeq \vec{v}$, $t \preccurlyeq \vec{w}$,
- there is some i with $u'_i = (v_i w_i lk)$, and $u'_j = v_j = w_j$ for $j \neq i$.

It is immediate to check that $\preccurlyeq'$ is a simulation relation. First, (old) states from S can only be simulated by (old) states from Q. Second, a new state $(stlj)$ of $\mathcal{B}$ can be simulated only by states $\vec{u}' \in Q' \setminus Q$. It can be shown easily that the largest simulation relation from $\mathcal{B}'$ to $\mathcal{A}'_1 \otimes \cdots \otimes \mathcal{A}'_n$ coincides with $\preccurlyeq'$ (hence with $\preccurlyeq$) on the set $S \times Q$ of pairs of old states.

□

5. The complexity of bisimulation

Till now we wanted to decide if an asynchronous product of deterministic automata $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ can simulate a deterministic automaton $\mathcal{B}$. An evident question is to consider what happens if we consider bisimulation instead of simulation. To be bisimilar to an asynchronous product, $\mathcal{B}$ must satisfy some structural constraints. In this section we prove the following theorem, which shows that indeed, the bisimulation problem is easier.

Theorem 5.1. *The following question can be solved in logarithmic space:*
   Input: *n deterministic automata $\mathcal{A}_1, \ldots, \mathcal{A}_n$ and a deterministic automaton $\mathcal{B}$.*
   Output: *decide if $\mathcal{B}$ and $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ are bisimilar.*

The proof of the theorem will occupy the rest of the section. We fix $\mathcal{B}$ and $\mathcal{A}_1, \ldots, \mathcal{A}_n$. Without loss of generality we assume that $\mathcal{B}$ is minimal with respect to bisimulation: no two different states of $\mathcal{B}$ are bisimilar (if $\mathcal{B}$ is not minimal we can minimize it on-the-fly in logarithmic space). This assumption also has a very pleasant consequence. If two states s_1 and s_2 of $\mathcal{B}$ are bisimilar to the same global state of $\mathcal{A}$, then $s_1 = s_2$.

As we aim to obtain a logarithmic space algorithm we cannot even allow ourselves to explore the state space of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ at random, as we cannot store the tuples of states. This is why the following definition is crucial for the construction.

Definition 5.2. A sequence of transitions of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ is *banal* if it can be decomposed into a, possibly empty, sequence of transitions of $\mathcal{A}_1$, followed by one of $\mathcal{A}_2$, and so on, up to $\mathcal{A}_n$.

Observe that thanks to the lack of synchronization every state of $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ is reachable by a run that is a banal sequence. Another pleasant property is that banal sequences can be explored in logarithmic space: we need only to remember the current state of the unique process that is active. We call *configuration* a pair $(s, \vec{t})$ consisting of a state s of $\mathcal{B}$ and a global state $\vec{t}$ of $\mathcal{A}$. For convenience, we say that a configuration $(s, \vec{t})$ is reachable by some sequence ρ of transitions of $\mathcal{A}$ if ρ leads to $\vec{t}$ from the initial state of $\mathcal{A}$, and if s is reached in $\mathcal{B}$ from the initial state by the sequence of actions associated with ρ (this is well-defined since $\mathcal{B}$ is deterministic). Note also that we can explore any configuration $(s, \vec{t})$ that is reachable by some banal sequence in logarithmic space. Let us call such pairs *banally-reachable configurations*.

The first necessary condition for $\mathcal{B}$ being bisimilar to $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ is that for every banally-reachable configuration $(s, \vec{t})$ the same actions are possible from s and $\vec{t}$. This can be checked in logarithmic space as it is easy to verify its negation within this bound.

The second necessary condition is that every reachable configuration is banally-reachable. Indeed, if $(s, \vec{t})$ is reachable by a sequence that is not banal then the banal sequence ρ obtained by ordering the transitions process-wise also reaches $\vec{t}$. If a bisimulation exists then we are guaranteed that ρ reaches s in $\mathcal{B}$. This is because the state reached by ρ must be bisimilar to s, and $\mathcal{B}$ is minimal with respect to bisimulation.

To show that one can check in logarithmic space that every reachable configuration is banally-reachable, we consider the negation of this property. We can then use the fact that LOGSPACE is closed under complement. We want to find a reachable configuration that is not banally-reachable. If one exists then we can look at one that is reachable in a shortest

number of steps. This means that there must exist a banally-reachable configuration $(s_1, \vec{t_1})$, an action b and a process i such that $(s_2, \vec{t_2})$ is not banally-reachable, where $\delta_{\mathcal{B}}(s_1, b) = s_2$ and $\vec{t_2}$ is obtained from $\vec{t_1}$ by taking transition b of process i. This can be checked as follows. One produces on-the-fly a banal sequence, when the part of process i is finished an extra transition with letter b is taken. This way we have two states, one before taking b and one after. We then continue constructing banal sequences from the two states with transitions of processes $i + 1$ up to n. This way we have obtained two sequences which differ by the action b of process i, and we check that the two states reached by $\mathcal{B}$ are different.

Together, the two conditions above are also sufficient for $\mathcal{A}_1 \otimes \cdots \otimes \mathcal{A}_n$ and $\mathcal{B}$ being bisimilar, hence the result.

6. Conclusion

We have shown an Exptime lower bound for the composition of services that are described as a fully asynchronous product of finite state machines. Thus, we answer the question left open in [2]. Since our lower bound holds for the simplest parallel composition operation one can think of (no synchronization at all), it also applies to richer models, such as products with synchronization on actions as in [10] or communicating finite-state machines (CFSM) as in [9, 8]. It is easy to see that the simulation of a finite-state machine by a CFSM $\mathcal{A}$ with bounded message queues is in Exptime, since the state space of $\mathcal{A}$ is exponential in this case. Hence, this problem, as well as any of its variants with some restricted form of communication, is Exptime-complete as well.

An interesting open question is what happens if we allow in the asynchronous product arbitrary many copies of each finite state machine. That is, we suppose that an available service can be used by an arbitrary number of peers. This question reduces to a bounded variant of the simulation of a finite state machine by a BPP, and its decidability status is open.

Acknowledgement: We thank the anonymous referees for interesting comments and suggestions for improvement.

References

[1] G. Alonso, F. Casati, H. Kuno, and V. Machiraju. *Web Services. Concepts, Architectures and Applications.* Springer, 2004.

[2] D. Berardi, D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Mecella. Automatic composition of e-services that export their behavior. In *Proc. of the 1st Int. Conf. on Service Oriented Computing (ICSOC 2003)*, LNCS 2910, pp. 43–58, 2003.

[3] D. Berardi, D. Calvanese, G. D. Giacomo, R. Hull, and M. Mecella. Automatic composition of web services in Colombo. In *SEBD 2005*, pages 8–15, 2005.

[4] D. Berardi, D. Calvanese, G. D. Giacomo, R. Hull, and M. Mecella. Automatic composition of web services with messaging. In *VLDB 2005*, pages 613–624, 2005.

[5] D. Brand and P. Zafiropulo. On communicating finite-state machines. *J. ACM* 30(2):323–342, 1983.

[6] A. K. Chandra, D. Kozen and L. J. Stockmeyer. Alternation. *J. ACM* 28(1):114–133, 1981.

[7] A. Deutsch, L. Sui, V. Vianu, and D. Zhou. Verification of communicating data-driven web services. In *Symposium on Principles of Database Systems (PODS)*, pp. 90-99, 2006.

[8] X. Fu, T. Bultan, and J. Su. Conversation protocols: a formalism for specification and verification of reactive electronic services. In *Theor. Comput. Sci.* 328(1-2):19–37, 2004.

[9] R. Hull, M. Benedikt, V. Christophides, J. Su. E-services: a look behind the curtain. In *Symposium on Principles of Database Systems (PODS)*, pp. 1-14, 2003.

[10] F. Laroussinie and Ph. Schnoebelen. The state explosion problem from trace to bisimulation equivalence. In *FoSSaCS 2000*, LNCS 1784, pp. 192–207, 2000.

[11] A. Muscholl and I. Walukiewicz. A lower bound on Web services composition. In *Proc. of FoSSaCS'07*, LNCS 4423, pp. 274-286, 2007.

[12] J. Srba. Roadmap of infinite results. *Bulletin of the EATCS* 78, pages 163-175, 2002. See also `http://www.brics.dk/∼srba/roadmap`.