

LOWERBOUNDS FOR BISIMULATION BY PARTITION REFINEMENT

JAN FRISO GROOTE , JAN MARTENS , AND ERIK P. DE VINK 

Eindhoven University of Technology, The Netherlands
e-mail address: j.f.groote@tue.nl, j.j.m.martens@tue.nl, e.p.d.vink@tue.nl

ABSTRACT. We provide time lowerbounds for sequential and parallel algorithms deciding bisimulation on labelled transition systems that use partition refinement. For sequential algorithms this is $\Omega((m+n)\log n)$ and for parallel algorithms this is $\Omega(n)$, where n is the number of states and m is the number of transitions. The lowerbounds are obtained by analysing families of deterministic transition systems, ultimately with two actions in the sequential case, and one action for parallel algorithms.

For deterministic transition systems with one action, bisimilarity can be decided sequentially with fundamentally different techniques than partition refinement. In particular, Paige, Tarjan, and Bonic give a linear algorithm for this specific situation. We show, exploiting the concept of an oracle, that this approach is not of help to develop a faster generic algorithm for deciding bisimilarity. For parallel algorithms there is a similar situation where our approach can be applied, too.

1. INTRODUCTION

Strong bisimulation [Par81, Mil80] is the gold standard for equivalence on labelled transition systems (LTSs). Deciding bisimulation equivalence among the states of an LTS is a crucial step for tool-supported analysis and model checking of LTSs. The well-known and widely-used partition refinement algorithm of Paige and Tarjan [PT87] has a worst-case upperbound $O(m\log n)$ for establishing the bisimulation equivalence classes. Here, n is the number of states and m is the number of transitions in an LTS.

The algorithm of Paige and Tarjan seeks to find, starting from an initial partition, via refinement steps, the coarsest stable partition, that in fact is built from the bisimulation equivalence classes that are looked for. The algorithm achieves the complexity of the logarithm of the number of states times the number of transitions by restricting the amount of work for refining blocks and moving states. When refining, the splitting blocks are investigated using an intricate bookkeeping trick. Only the smaller parts of a block that are to be moved to a new block are split off, leaving the bulk of the original block at its place. These specific ideas go back to [Hop71] and make the difference with the earlier $O(mn)$ algorithm of Kanellakis and Smolka [KS90]. The algorithms by Kanellakis-Smolka and Paige-Tarjan, with the format of successive refinements of an initial partition till a

Key words and phrases: Bisimilarity, partition refinement, labelled transition system, lowerbound.
Partially funded by the AVVA project NWO 612.001.751/TOP1.17.002.

fixpoint is reached, have been leading for variations and generalisations for deciding specific forms of (strong) bisimilarities, see e.g. [Buc99, DPP04, GVV18, WDMS20, JGKW20].

We are interested in the question whether the Paige-Tarjan algorithm is computationally optimal. A lowerbound for a related problem is provided in [BBG17] that studies colour refinement of graphs. Colour refinement computes, given a graph and an initial colouring, a minimal consistent colouring such that every two equally coloured nodes have, for every colour, the same number of neighbours of the same colour. More specifically, in that paper it is proven that for a family of graphs with n nodes and m edges, finding the canonical coarsest stable colouring is in $\Omega((m+n)\log n)$. However, the costs for computations on graphs for colour refinement are charged differently than those for partition refinement for bisimulation on LTSs. The former takes edges between blocks of uniformly coloured nodes into account, the latter focuses on the size of newly created blocks of states. In [BBG17] it is described how the family of graphs underlying the lowerbound for colour refinement can be transformed into a family of Kripke structures for which computing bisimulation is $\Omega((m+n)\log n)$ when counting the numbers of edges.

In this paper we follow a different approach to obtain a lowerbound. We define the concept of a partition refinement algorithm and articulate the complexity in terms of the number of states that are moved. In particular, we define the notion of a valid refinement sequence which has its counterpart in iteration sequences for colour refinement. Then, we introduce a family of (deterministic) LTSs, called bisplitters, for which we show that computing bisimulation requires $n \log n$ work. The family of $n \log n$ -hard LTSs that we use to establish the lowerbound, involves an action set of $\log n$ actions. Building on this result and exploiting ideas borrowed from [PTB85] to extend the bisimulation classes for the states in the end structures, i.e. cycles, to the states of the complete LTS, we provide another family of (deterministic) LTSs that have two actions only. Then we argue that for the two-action case the complexity of deciding bisimulation is $\Omega((m+n)\log n)$. We want to stress that the families involved consist of deterministic LTSs.

Recently, a linear time algorithm for bisimilarity was proposed for a PRAM (Parallel Random Access Machine) using $\max(n, m)$ processors [MGH⁺21]. This algorithm also employs partition refinement. This naturally raises the question whether the algorithm is optimal, or whether it can fundamentally be improved. We answer the question in the present paper by showing an $\Omega(n)$ lowerbound for parallel algorithms employing partition refinement, using a family of deterministic transition systems with one action label.

We obtain our lowerbound results assuming that algorithms use partition refinement. However, one may wonder if a different approach than partition refinement can lead to a faster decision procedure for bisimulation. For the specific case of deterministic LTSs with a singleton action set and state labelling, Robert Paige, Robert Tarjan and Robert Bonic propose a sequential algorithm [PTB85] that uses linear time. We refer to it as Roberts' algorithm. In [CRS08] it is proven that partition refinement à la Hopcroft has a lowerbound of $\Omega(n \log n)$ in this case. Concretely, this means that in the one-letter case Roberts' algorithm achieves the essentially better performance by using a completely different technique than partition refinement to determine the bisimulation equivalence classes.

Crucial for Roberts' algorithm is the ability to identify, in linear time, the bisimilarity classes of cycles. In this paper we show that if the alphabet consists of at least two actions, a rapid decision on 'cycles' as in [PTB85] will not be of help to improve on the Paige-Tarjan algorithm for general LTSs. We argue that the speciality in the algorithm of [PTB85], viz. to be able to quickly decide the bisimilarity of the states on a cycle, can be captured by

means of a stronger notion, namely an oracle that provides the bisimulation classes of the states of a so-called ‘end structure’, the counterpart in the multiple action setting of a cycle in the single action setting. The oracle can be consulted to refine the initial partition with respect to the bisimilarity on the end structures of the LTS for free. We show that for the class of sequential partition refinement algorithms enhanced with an oracle as described, thus encompassing the algorithm of [PTB85], the $(m+n) \log n$ lowerbound persists for action sets with at least two actions.

For parallel algorithms a similar situation occurs as for deterministic Kripke structures: an $O(\log n)$ parallel algorithm exists [JR94] to determine the bisimulation equivalence classes. This algorithm also necessarily employs techniques that go beyond partition refinement. We believe that these techniques cannot be used either to fundamentally improve the complexity of determining bisimilarity on LTSs, but leave the proof as an open question.

The document is structured as follows. In Section 2 we give the necessary preliminaries on the problem. A recap of the linear algorithm of [PTB85] is provided in Section 3. Next, we introduce the family of LTSs \mathcal{B}_k for which we show in Section 4 that deciding bisimilarity is $\Omega(n \log n)$ for the class of partition refinement algorithms and for which we establish in Section 5 an $\Omega(n \log n)$ lowerbound for the class of partition refinement algorithms enhanced with an oracle for end structures. In Section 6 we introduce the family of deterministic LTSs \mathcal{C}_k , each involving two actions only, to take the number of transitions m into account and establish an $\Omega((m+n) \log n)$ lowerbound for partition refinement with and without an oracle for end structures. In Section 7 we provide the $\Omega(n)$ lowerbound for parallel refinement algorithms. In Section 8 we discuss the differences and similarities with the lowerbound results on colour refinement of [BBG17]. We wrap up with concluding remarks.

Note The present paper is an extension the conference paper [GMV21] that appeared in the proceedings of CONCUR 2021.

2. PRELIMINARIES

Given a set of states S , a *partition* of S is a set of subsets of states $\pi \subseteq 2^S$ such that $\emptyset \notin \pi$, for all $B, B' \in \pi$ either $B \cap B' = \emptyset$ or $B = B'$, and $\bigcup_{B \in \pi} B = S$. The elements of a partition are referred to as blocks. A partition π of S induces an equivalence relation $=_\pi \subseteq S \times S$, where for two states $s, t \in S$, $s =_\pi t$ iff the states s and t are in the same block, i.e. there is a block $B \in \pi$ such that $s, t \in B$. A partition π' of S is a *refinement* of a partition π of S iff for every block $B' \in \pi'$ there is a block $B \in \pi$ such that $B' \subseteq B$. It follows that each block of π is the disjoint union of blocks of π' . The refinement is *strict* if $\pi \neq \pi'$. The common refinement of two partitions π and π' is the partition with blocks $\{B \cap B' \mid B \in \pi, B' \in \pi': B \cap B' \neq \emptyset\}$. A sequence of partitions (π_0, \dots, π_n) is called a refinement sequence iff π_{i+1} is a refinement of π_i , for all $0 \leq i < n$.

Definition 2.1. A labelled transition system with initial partition (LTS) is a four-tuple $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ where S is a finite set of states S , \mathcal{A} is a finite alphabet of actions, $\rightarrow \subseteq S \times \mathcal{A} \times S$ is a transition relation, and π_0 is a partition of the set of states S . A labelled transition system with initial partition is called deterministic if the transition relation is a total function $S \times \mathcal{A} \rightarrow S$.

Given an LTS $L = (S, \mathcal{A}, \rightarrow, \pi_0)$, states $s, t \in S$, and an action $a \in \mathcal{A}$, we write $s \xrightarrow{a} t$ instead of $(s, a, t) \in \rightarrow$. For notational convenience, we occasionally write $L[U] = \{t \in S \mid$

$\exists s \in U \exists a \in \mathcal{A}: s \xrightarrow{a} t \}$, and, for a deterministic LTS L , we may use $L(s, a)$ to denote the unique state t of L such that $s \xrightarrow{a} t$. We say that s reaches t via a iff $s \xrightarrow{a} t$. A state s reaches a set $U \subseteq S$ via action a iff there is a state in U that is reached by s via a , notation $s \xrightarrow{a} U$. A set of states $V \subseteq S$ is called *stable* under a set of states $U \subseteq S$ iff for all actions a , either all states in V reach U via a , or no state in V reaches U via a . Thus, a set of states V is *not* stable under U iff for two states s and t in V and an action a it holds that $s \xrightarrow{a} U$ and $t \not\xrightarrow{a} U$. A partition π is stable under a set of states U iff each block $B \in \pi$ is stable under U . A partition π is called *stable* iff it is stable under all its blocks. So, for any two blocks B and C of π and any action $a \in \mathcal{A}$, either each state s of B has an a -transition to C or each state s of B doesn't have an a -transition to C .

Following [Par81, Mil80], given an LTS L , a symmetric relation $R \subseteq S \times S$ is called a bisimulation relation iff for all $(s, t) \in R$ and $a \in \mathcal{A}$, we have that $s \xrightarrow{a} s'$ for some $s' \in S$ implies that $t \xrightarrow{a} t'$ for some $t' \in S$ such that $(s', t') \in R$. In the setting of the present paper, as we incorporate the initial partition in the definition of an LTS, bisimilarity is slightly non-standard. For a bisimulation relation R , we additionally require that it respects the initial partition π_0 of L , i.e. $(s, t) \in R$ implies $s =_{\pi_0} t$. Two states $s, t \in S$ are called (strongly) bisimilar for L iff a bisimulation relation R exists with $(s, t) \in R$, notation $s \leftrightarrow_L t$. Bisimilarity is an equivalence relation on the set of states of L . We write $[s]_L^{\leftrightarrow}$ for the bisimulation equivalence class of the state s in L .

Note that for a deterministic LTS with a set of states S and initial partition $\pi_0 = \{S\}$, we have that π_0 itself already represents bisimilarity, contrary to LTSs in general.

Partition refinement algorithms for deciding bisimilarity on LTSs start with an initial partition π_0 , which is subsequently repeatedly refined until a stable partition is reached. Thus, unstable blocks are replaced by several smaller blocks. The stable partition that is reached happens to be the coarsest stable partition of the LTS refining π_0 and coincides with bisimilarity [KS90, PT87].

Below we define so-called *valid* refinement sequences. An algorithm is called a partition refinement algorithm iff every run of the algorithm is reflected by a valid refinement sequence (π_0, \dots, π_n) . All the lowerbounds that we provide apply to algorithms producing valid partition sequences, which virtually all known bisimulation algorithms do, and as such this is the core definition in this paper.

A partition sequence (π_0, \dots, π_n) is valid when the direct successor π_i of a partition π_{i-1} in the sequence is obtained by splitting one or more unstable blocks in π_{i-1} using only information available in π_{i-1} . Furthermore, the last partition in the sequence, the partition π_n , is stable. If block B of π_{i-1} is replaced in π_i because it is not stable under block B' of π_{i-1} , then B' is referred to as a splitter block.

Definition 2.2. Let $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ be an LTS, and π a partition of S . A refinement π' of π is called a *valid refinement* with respect to L iff the following criteria hold.

- (a) π' is a strict refinement of π .
- (b) If $s \neq_{\pi'} t$ for $s, t \in S$, then (i) $s \neq_{\pi} t$ or (ii) $s' \in S$ exists such that $s \xrightarrow{a} s'$ for some $a \in \mathcal{A}$ and, for all $t' \in S$ such that $t \xrightarrow{a} t'$, it holds that $s' \neq_{\pi} t'$, or the other way around with t replacing s .

A sequence of partitions $\Pi = (\pi_0, \dots, \pi_n)$ is called a valid partition sequence iff every successive partition π_i , for $0 < i \leq n$, is a valid refinement of π_{i-1} , and, moreover, the partition π_n is stable.

When a partition π is refined into a partition π' , states that are in the same block but can reach different blocks can lead to a split of the block into smaller subsets, say k subsets. This means that a block $B \in \pi$ is split into blocks $B_1, \dots, B_k \in \pi'$. The least amount of work is done for this operation if we create new blocks for the least number of states. That means if $B \in \pi$ is split into $B_1, \dots, B_k \in \pi'$ and B_1 is the biggest block, then the states of B_2, \dots, B_k are moved to new blocks and the states of B_1 remain in the current block that was holding B . Therefore, we define the refinement costs rc for the refinement π' of π by

$$rc(\pi, \pi') = \sum_{B \in \pi} (|B| - \max_{B' \in \pi': B' \subseteq B} |B'|).$$

For a sequence of refinements $\Pi = (\pi_0, \dots, \pi_n)$ we write $rc(\Pi)$ for $\sum_{i=1}^n rc(\pi_{i-1}, \pi_i)$. For an LTS L , we have

$$rc(L) = \min\{rc(\Pi) \mid \Pi \text{ a valid refinement sequence for } L\}.$$

Note that this complexity measure is different from the one used in [BBG17], which counts transitions. Our complexity measure rc is bounded from above by the former.

In various examples below we characterise the states of LTSs by sequences of bits. The set of bits is denoted as $\mathbb{B} = \{0, 1\}$. Bit sequences of length up to and including k are written as $\mathbb{B}^{\leq k}$. The complement of a bit b is denoted by \bar{b} . Thus $\bar{0} = 1$ and $\bar{1} = 0$. For two bit sequences σ, σ' , we write $\sigma \preceq \sigma'$ to indicate that σ is a prefix of σ' and write $\sigma \prec \sigma'$ iff σ is a strict prefix of σ' . For a bit sequence $\sigma \in \mathbb{B}^k$, for any $i, j \leq k$, we write $\sigma[i]$ to indicate the bit at position i starting from position 1. We write $\sigma[i:j] = \sigma[i]\sigma[i+1] \cdots \sigma[j]$ to indicate the subword from position i to position j . Occasionally we use, for a bit sequence σ , the notation $\sigma\mathbb{B}^k$ to denote $\{\sigma\sigma' \mid \sigma' \in \mathbb{B}^k\}$, the set of all bit sequences of length $|\sigma| + k$ having σ as prefix.

3. ROBERTS' ALGORITHM

Most algorithms to determine bisimulation for an LTS use partition refinement. However, there are a few notable exceptions to this. For the class of deterministic LTSs that have a singleton action alphabet, deciding the coarsest stable partition, i.e. bisimilarity, requires linear time only; a linear algorithm is due to Robert Paige, Robert Tarjan, and Robert Bonic [PTB85], which we therefore aptly refer to as Roberts' algorithm.

The algorithm of [PTB85] exploits the specific structure of a deterministic LTS with one action label. An example of such a transition system is depicted in Figure 1, where the action label itself has been suppressed and the initial partition is indicated by single/double circled states. In general, a deterministic LTS with one action label can be characterised as a directed graph, possibly with self-loops, consisting of a number of cycles of one or more states together with root-directed trees with their root on a cycle. Below we refer to a cycle with the trees connected to it as an end structure. In a deterministic LTS with one action label, each state belongs to a unique end structure; it is on a cycle or has a unique directed path leading to a cycle.

In brief, Roberts' algorithm for deterministic LTSs with one action label can be described as follows (see [PTB85] for more details).

- (1) As preparatory step, find all the end structures of the LTS, i.e. detect all cycles, and all root-oriented trees leading to cycles.

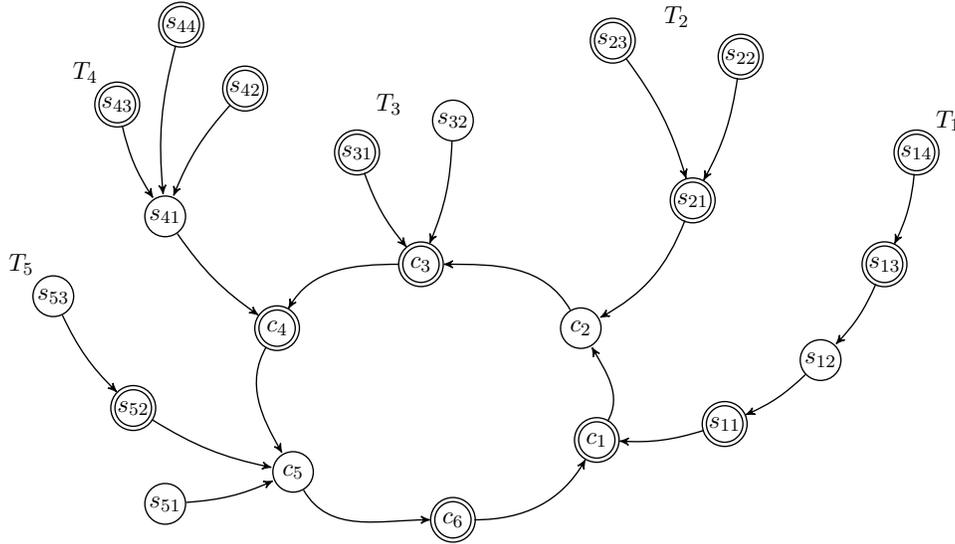


Figure 1: An example of a deterministic LTS with initial partition (action label suppressed).

- (2) Observe that each state s on a cycle encodes a sequence of blocks, viz. the sequence starting from the block the state is in, and blocks encountered when following the transitions, up to the state on the cycle that leads back to s . This sequence of blocks forms a word w over the alphabet of the initial partition, where each block of the initial partition is a symbol of this alphabet. The word w can be uniquely written as v^k with v of minimal length and $k > 0$. The string v is referred to as the repeating prefix of the state s .

We consider the repeating prefixes of all states on the cycle and identify the lexicographically least repeating prefix v . This can be done in linear time in the size of the cycle using a string matching algorithm due to Knuth, Morris, and Pratt [KMP77]. The lexicographically least repeating prefix v and the minimal number of transitions that is required to reach a state t from a state that has v as repeating prefix, determines the bisimulation equivalence class of the state t . We encode this bisimulation equivalence class by the corresponding rotation of the prefix v . This way the bisimulation class is established for all states on all cycles. By comparing least repeating prefixes bisimilarity across cycles can be detected.

- (3) By a backward calculation along the path leading from a state up in a tree down to their root on a cycle, the bisimilarity equivalence classes for the remaining states can subsequently be determined in linear time as well. The root of the tree is a state on the cycle and therefore has been assigned a string, hence a bisimulation class. We assign to a child the string of the parent prepended with the symbol of the initial class of the child.

Example. The deterministic LTS of Figure 1 has a single end structure, viz. the cycle formed by the states c_1 to c_6 and five trees, tree T_1 with leaf s_{14} and rooted in c_1 , tree T_2 with leaves s_{22} and s_{23} rooted in c_2 , tree T_3 with leaves s_{31} and s_{32} rooted in c_3 , the tree T_4

with leaves s_{42} , s_{43} , and s_{44} rooted in c_4 , and the tree T_5 with leaves s_{51} and s_{53} rooted in c_5 .

With the symbol A for an accepting, i.e. double-circled, state and the symbol N for a non-accepting state, i.e. not double-circled, we associate the following sequences of A 's and N 's to the nodes on the cycle:

$$\begin{array}{lll} c_1 : ANAANA & c_3 : AANAAN & c_5 : NAANAA \\ c_2 : NAANAA & c_4 : ANAANA & c_6 : AANAAN. \end{array}$$

With A preceding N , the lexicographically least repeating prefix is AAN . We assign states c_1 and c_4 to the bisimulation class of ANA , states c_2 and c_5 to the bisimulation class of NAA , and states c_3 and c_6 to the bisimulation class of AAN . Here, ANA and NAA are the 1-place and 2-place rotations of AAN , respectively.

Moving to tree T_1 with root c_1 having string ANA associated with it, we concatenate for state s_{11} the symbol A , since s_{11} has been assigned in the block of accepting states initially, followed by the string ANA of c_1 , forming $AANA$ which is reduced to AAN (exploiting the equality $A(ANA)^\omega = (AAN)^\omega$). Thus, we see that the states s_{11} and c_6 are bisimilar. Similarly, for state s_{12} we prepend the symbol N of the child s_{12} to the string AAN of the parent s_{11} and obtain $NAAN$ which is reduced to NAA (now exploiting the equality $N(AAN)^\omega = (NAA)^\omega$), as for state c_5 . Considering in contrast the state s_{41} , with symbol N and which is child of state c_4 with string ANA , s_{41} gets assigned the string $NANA$ (which can not be reduced, since $N(ANA)^\omega \neq (NAN)^\omega$). Subsequently, the states s_{42} to s_{44} get assigned $ANANA$.

Roberts' algorithm solves the so-called single function coarsest partition problem in $O(n)$ for a set of n elements. A striking result is that any algorithm that is based on partition refinement requires $\Omega(n \log n)$, as witnessed in [BC04, CRS08], where it is shown that partition refinement algorithm of Hopcroft [Hop71] cannot do better than $O(n \log n)$. Thus, Roberts' algorithm must use other techniques than partition refinement. Below we come back to this observation, showing that it is not possible to use the ideas in Roberts' algorithm to come up with a linear algorithm for computing bisimilarity for a class of LTSs that either includes nondeterministic LTSs, or allows LTSs to involve more than one action label.

4. \mathcal{B}_k IS $\Omega(n \log n)$ FOR PARTITION REFINEMENT

In this section we introduce a family of deterministic LTSs called bisplitters \mathcal{B}_k for $k \geq 1$, on which the cost of any partition refinement algorithm is $\Omega(n \log n)$, where n is the number of states. Building on the family of \mathcal{B}_k 's, we propose in Section 6 a family of LTSs \mathcal{C}_k for which the cost of partition refinement is $\Omega((n+m) \log n)$, where m is the number of transitions.

Definition 4.1. For $k \geq 1$, the bisplitter \mathcal{B}_k is defined as the LTS with initial partition $\mathcal{B}_k = (\mathbb{B}^k, \mathcal{A}_k, \rightarrow, \pi_0^k)$ where the set of states \mathbb{B}^k is the set of all bit strings of length k , $\mathcal{A}_k = \{a_1, \dots, a_{k-1}\}$ is a set of $k-1$ actions, the transition relation is given by

$$\begin{aligned} & \{ \sigma \xrightarrow{a_i} \sigma \mid \sigma \in \mathbb{B}^k, 1 \leq i < k: \sigma[i+1] = 0 \} \cup \\ & \{ \sigma \xrightarrow{a_i} \sigma[1:i-1] \overline{\sigma[i]} 0^{k-i} \mid \sigma \in \mathbb{B}^k, 1 \leq i < k: \sigma[i+1] = 1 \}, \end{aligned}$$

and $\pi_0^k = \{B_0, B_1\}$, where $B_0 = 0\mathbb{B}^{k-1}$ and $B_1 = 1\mathbb{B}^{k-1}$, is the initial partition.

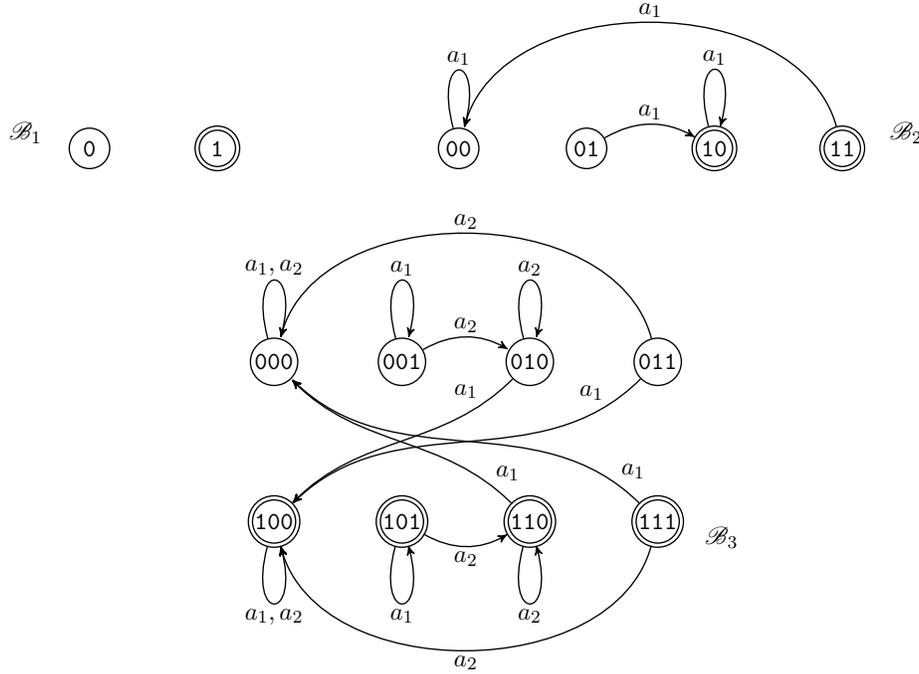


Figure 2: The bisplitters \mathcal{B}_1 , \mathcal{B}_2 , and \mathcal{B}_3 . Initial partitions are indicated by single-circled and double circled states.

We see that the bisplitter \mathcal{B}_k has 2^k states, viz. all bit strings of length k , and $k-1$ different actions. The LTS \mathcal{B}_k is deterministic. Each state has exactly one outgoing transition for each action a_i , $1 \leq i < k$. Thus, \mathcal{B}_k has $(k-1)2^k$ transitions: (i) a self-loop for bitstring σ with label a_i if the $i+1$ -th bit $\sigma[i+1]$ of σ equals 0; (ii) otherwise, i.e. when bit $\sigma[i+1]$ equals 1, the bitstring σ has a transition for label a_i to the bitstring that equals the first $i-1$ bits of σ , flips the i -th bit of σ , and has $k-i$ many 0's following. The initial partition π_0^k distinguishes the bit strings starting with 0 from those starting with 1.

Drawings of the first three bisplitters \mathcal{B}_1 to \mathcal{B}_3 are given in Figure 2. We see in the picture of \mathcal{B}_3 for example for the bitstring $\sigma = 101$, an a_1 -transition to itself, as $\sigma[1+1] = \sigma[2] = 0$ and an a_2 -transition to the bitstring 110, as $\sigma[2+1] = \sigma[3] = 1$, $\overline{\sigma[2]} = \overline{0} = 1$, and $\sigma[1]\overline{\sigma[2]}0 = 110$. As another illustration, for the bitstring $\sigma = 011$ we have an a_1 -transition to $\sigma[1]00 = 100$ since $\sigma[1+1] = 1$ and an a_2 -transition to $\sigma[1]\overline{\sigma[2]}0 = 000$ since $\sigma[2+1] = 1$.

Also note that \mathcal{B}_3 contains two copies of \mathcal{B}_2 . In the copies, the action label a_1 of \mathcal{B}_2 maps to the action label a_2 in \mathcal{B}_3 , and each state associated with a bitstring $\sigma \in \mathbb{B}^2$ produces two copies in \mathcal{B}_3 ; one copy is obtained by the mapping $\sigma \mapsto 0\sigma$ and the other copy is obtained by the mapping $\sigma \mapsto 1\sigma$. In general, bisplitter \mathcal{B}_k is twice embedded in bisplitter \mathcal{B}_{k+1} via the mappings $\sigma \mapsto 0\sigma$ and $\sigma \mapsto 1\sigma$ from \mathbb{B}^k to \mathbb{B}^{k+1} for the states using the mapping $a_i \mapsto a_{i+1}$ from \mathcal{A}_k to \mathcal{A}_{k+1} for the action labels. Note that initial partitions are not respected.

Definition 4.2. For any string $\sigma \in \mathbb{B}^{\leq k}$, we define the prefix block B_σ of \mathcal{B}_k to be the block $B_\sigma = \{\sigma' \in \mathbb{B}^k \mid \sigma \preceq \sigma'\}$.

The following lemma collects a number of results related to prefix blocks that we need in our complexity analysis for computing bisimilarity for the bisplitters.

Lemma 4.3. *Let $k \geq 1$ and consider the LTS with initial partition $\mathcal{B}_k = (\mathbb{B}^k, \mathcal{A}_k, \rightarrow, \pi_0^k)$, i.e. the k -th bisplitter. Let the sequence $\Pi = (\pi_0^k, \dots, \pi_n)$ be a valid refinement sequence for \mathcal{B}_k . Then it holds that*

- (a) *Every partition π_i in Π contains prefix blocks only.*
- (b) *If partition π_i in Π contains a prefix block B_σ with $|\sigma| < k$, then π_i is not stable.*
- (c) *If B_σ is in π_i , for $0 \leq i < n$, then either $B_\sigma \in \pi_{i+1}$, or $B_{\sigma_1} \in \pi_{i+1}$ and $B_{\sigma_0} \in \pi_{i+1}$.*

Proof. (a) Initially, for $\pi_0^k = \{B_0, B_1\}$, both its blocks are prefix blocks by definition. We prove, if partition π_i , for $0 \leq i < n$, consists of prefix blocks only, then all blocks in π_{i+1} are prefix blocks as well.

Assume, to arrive at a contradiction, that there is a block $B \in \pi_{i+1}$ that is not a prefix block. Because π_{i+1} is a refinement of π_i , we have $B \subseteq B_\sigma$ for some prefix block $B_\sigma \in \pi_i$. This means that σ is a common prefix of all elements of B . We can choose θ such that $\sigma\theta$ is the longest common prefix of all elements of B . Since every singleton of \mathbb{B}^k is a prefix block, B is not a singleton. This implies that $|\sigma\theta| < k$ and that there are elements σ_1 and σ_2 of B such that $\sigma\theta\sigma_1$ is a prefix of σ_1 and $\sigma\theta\sigma_2$ is a prefix of σ_2 . Because B is not a prefix block by assumption, there must exist a string $\tau \in \mathbb{B}^k$ with prefix $\sigma\theta$ such that $\tau \notin B$. Obviously, we have either (i) $\sigma\theta\sigma_1$ is a prefix of τ , or (ii) $\sigma\theta\sigma_2$ is a prefix of τ . We will show that in both these cases τ in fact belongs to B , thus arriving at a contradiction.

- (i) Suppose $\sigma\theta\sigma_1$ is a prefix of τ . We argue that τ and σ_1 belong to the same block in π_{i+1} since, for each a_j , $1 \leq j < k$, the target states σ'_1 and τ' of the transitions $\sigma_1 \xrightarrow{a_j} \sigma'_1$ and $\tau \xrightarrow{a_j} \tau'$ belong to the same block of π_i . There are three cases:
 - $j < |\sigma\theta|$: Since $\sigma\theta$ is a prefix of both σ_1 and τ , we have $\sigma_1[j+1] = \tau[j+1]$.
 - If $\sigma_1[j+1] = \tau[j+1] = 0$, then $\sigma'_1 = \sigma_1$ and $\tau' = \tau$. Obviously, both σ'_1 and τ' belong to B_σ (since σ_1 and τ belong to B_σ).
 - If $\sigma_1[j+1] = \tau[j+1] = 1$, then both σ'_1 and τ' are of the form $\varrho[1:j-1]\overline{\varrho[j]}0^{k-j}$ where $\varrho = \sigma\theta$, and we have $\sigma'_1 = \tau'$. So, they clearly belong to the same block of π_i .
 - $j = |\sigma\theta|$: Since $\sigma_1[j+1] = \tau[j+1] = 0$, we have $\sigma'_1 = \sigma_1$ and $\tau' = \tau$, and hence both σ'_1 and τ' belong to B_σ .
 - $j > |\sigma\theta|$: In this case, for a string of the form $\sigma\theta\varrho$, an a_j -transition leads to a string of the form $\sigma\theta\varrho'$. In particular this means that if $j > |\sigma\theta|$ and $\sigma_1 \xrightarrow{a_j} \sigma'_1$ and $\tau \xrightarrow{a_j} \tau'$, then $\sigma\theta$ is a prefix of both σ'_1 and τ' , and σ'_1 and τ' belong to B_σ in π_i .
- (ii) Now, suppose $\sigma\theta\sigma_2$ is a prefix of τ . We argue that τ and σ_2 belong to the same block in π_{i+1} because for each a_j (where $1 \leq j < k$) the transitions $\sigma_2 \xrightarrow{a_j} \sigma'_2$ and $\tau \xrightarrow{a_j} \tau'$ lead to the same block of π_i . Also, here there are three cases:
 - $j < |\sigma\theta|$: Similar as for (i).
 - $j = |\sigma\theta|$: Since $\sigma_2[j+1] = \tau[j+1] = 1$, we have $\sigma'_2 = \tau' = \varrho[1:j-1]\overline{\varrho[j]}0^{k-1}$ where $\varrho = \sigma\theta$, so clearly σ'_2 and τ' are in a same block in π_i .
 - $j > |\sigma\theta|$: Similar as for (i).

Thus, both in case (i) and in case (ii) we see that we must have $\tau \in B$, contradicting the choice for τ .

(b) Suppose $B_\sigma \in \pi_i$ and $|\sigma| = \ell < k$. Let $\theta \in \mathbb{B}^*$ be such that $\sigma_1 = \sigma\theta$ and $\sigma_2 = \sigma\mathbf{1}\theta$. Then we have $\sigma_1 \xrightarrow{a_\ell} \sigma_1 \in B_\sigma$ and $\sigma_2 \xrightarrow{a_\ell} \sigma[1:\ell-1]\overline{\sigma[\ell]}0^{k-\ell} \notin B_\sigma$. Thus B_σ isn't stable, and hence π_i isn't either.

(c) We show that for a prefix block $B_\sigma \in \pi_i$, a bit $\mathbf{b} \in \mathbb{B}$ and all $\theta, \theta' \in \mathbb{B}^{k-(|\sigma|+1)}$, the states $\sigma_1 = \sigma\mathbf{b}\theta$ and $\sigma_2 = \sigma\mathbf{b}\theta'$ are not split by action a_j , for $1 \leq j < k$, and thus are in the same block of π_{i+1} . Pick j , $1 \leq j < k$, and suppose $\sigma_1 \xrightarrow{a_j} \sigma'_1$ and $\sigma_2 \xrightarrow{a_j} \sigma'_2$. If $j \leq |\sigma|$ and $\sigma[j+1] = 0$ then $\sigma'_1 = \sigma_1$ and $\sigma'_2 = \sigma_2$, hence both $\sigma'_1, \sigma'_2 \in B_\sigma$ don't split for a_j . If $j \leq |\sigma|$ and $\sigma[j+1] = 1$ then $\sigma'_1 = \sigma'_2$ and don't split for a_j either. If $j > |\sigma|$ then both $\sigma'_1, \sigma'_2 \in B_\sigma$ and don't split for a_j either. \square

With the help of the above lemma, clarifying the form of the partitions in a valid refinement sequence for the bisplitter family, we are able to obtain a lowerbound for any algorithm exploiting partition refinement to compute bisimilarity.

Theorem 4.4. *For any $k > 1$, application of partition refinement to the bisplitter \mathcal{B}_k has refinement costs $rc(\mathcal{B}_k) \in \Omega(n \log n)$ where $n = 2^k$ is the number of states of \mathcal{B}_k .*

Proof. Let $\Pi = (\pi_0^k, \dots, \pi_m)$ be a valid refinement sequence for \mathcal{B}_k . By items (a) and (b) of Lemma 4.3, we have $\pi_m = \{ \{s\} \mid s \in \mathbb{B}^k \}$ since π_m is stable and thus contains singleton blocks only. Item (c) of Lemma 4.3 implies that in every refinement step (π_i, π_{i+1}) a block is either kept or it is refined in two prefix blocks of equal size. The cost of refining the block B_σ , for $1 \leq |\sigma| \leq k-1$, into $B_{\sigma 0}$ and $B_{\sigma 1}$ is the number of states in $B_{\sigma 0}$ or the number of states in $B_{\sigma 1}$, which are the same and are equal to $\frac{1}{2}2^{k-|\sigma|}$. Therefore, we have

$$rc(\Pi) = \sum_{\ell=1}^{k-1} 2^\ell \frac{1}{2} 2^{k-\ell} = \sum_{\ell=1}^{k-1} \frac{1}{2} 2^k = (k-1)2^{k-1}.$$

With n the number of states of \mathcal{B}_k , we have that $n = 2^k$, thus $k-1 = \log \frac{1}{2}n$. Hence, $rc(\Pi) = \frac{1}{2}n \log \frac{1}{2}n$ which is in $\Omega(n \log n)$.

Thus, for every valid partition refinement sequence Π for \mathcal{B}_k we have $rc(\Pi) \in \Omega(n \log n)$. In particular this bound applies to the valid refinement sequence of minimal cost, and hence we conclude $rc(\mathcal{B}_k) \in \Omega(n \log n)$. \square

5. \mathcal{B}_k IS $\Omega(n \log n)$ FOR PARTITION REFINEMENT WITH AN ORACLE

In the previous section we have shown that computing bisimilarity with partition refinement for the family of bisplitters is $\Omega(n \log n)$. The bisplitters are deterministic LTSs but have growing actions sets. For the corner case of deterministic transition systems with a singleton action set, Roberts' algorithm discussed in Section 3 establishes bisimilarity in $O(n)$. Linearity was obtained by the trick of calculating the (lexicographically) least repeating prefix on the cycles in the transition system.

One may wonder whether an approach different from partition refinement of establishing bisimulation equivalence classes for transition systems with non-degenerate action sets can provide a linear performance. In order to capture the approach of [PTB85], we augment the class of partition refinement algorithms with an oracle. At the start of the algorithm the oracle can be consulted to identify the bisimulation classes for designated states, viz. for those that are in a so-called *end structure*, the counterpart of the cycles in Roberts'

algorithm. This results in a refinement of the initial partition; partition refinement then starts from the updated partition.

Thus, we can ask the oracle to provide the bisimulation classes of all states in an end structure of the input LTS, also including bisimilar states of the LTS not in an end structure. This yields a new partition, viz. the common refinement of the initial partition, on the one side, and the partition induced by the bisimulation equivalence classes as given by the oracle and the complement of their union, on the other side. Hence, the work that remains to be done is establishing the bisimulation equivalence classes, with respect to the initial partition, for the states not bisimilar to any in an end structure.

We will establish that a partition refinement algorithm that can consult an oracle cannot improve upon the complexity of computing bisimulation by partition refinement. We first define the notion of an end structure of an LTS formally as well as the associated notion of an end structure partition.

Definition 5.1. Given an LTS $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ with an initial partition, a non-empty subset $S' \subseteq S$ is called an *end structure* of L iff S' is a minimal set of states that is closed under all transitions, $L[S'] \subseteq S'$ and for all $S'' \subseteq S$ it holds that $L[S''] \subseteq S''$ and $S' \cap S'' \neq \emptyset$ implies $S' \subseteq S''$. Moreover, $es(L) = \{S' \subseteq S \mid S' \text{ end structure of } L\}$, $ES(L) = \bigcup es(L)$, and the partition π_{es} such that

$$\pi_{es} = \{[s]_L^{\leftrightarrow} \mid s \in ES(L)\} \cup \{B \setminus \bigcup_{s \in ES(L)} [s]_L^{\leftrightarrow} \mid B \in \pi_0\} \setminus \{\emptyset\}$$

is called the end structure partition of L .

Like the cycles exploited in Roberts' algorithm, an LTS can have multiple end structures. The end structure partition π_{es} consists of all the bisimilarity equivalence classes of L that include at least one state of an end structure, completed with blocks holding the remaining states, if non-empty. So, for every state s of an end structure, the end structure partition has identified all states that are bisimilar to state s and separates s and its bisimilar states from the rest of the LTS. The other states are assigned in the end structure partition to the blocks just as the initial partition does.

Example. In the LTS of Figure 1 the cycle of c_1 to c_6 is the only end structure. All states have a path to the cycle, hence every non-empty set that is closed under transitions will contain the cycle. Would the LTS have contained any isolated states, these would be end structures by themselves. Thus, consultation of the oracle leads to the refinement of the initial structure π_0 that consists of the two blocks

$$\{c_1, c_3, c_4, c_6\} \cup \{s_{11}, s_{13}, s_{14}\} \cup \{s_{21}, s_{22}, s_{23}\} \cup \{s_{31}\} \cup \{s_{42}, s_{43}, s_{44}\} \cup \{s_{52}\} \\ \text{and } \{c_2, c_5\} \cup \{s_{12}\} \cup \{s_{32}\} \cup \{s_{41}\} \cup \{s_{51}, s_{53}\}$$

into end structure partition π_{es} with five blocks, viz. the three blocks

$$\{c_1, c_4\} \cup \{s_{13}, s_{21}, s_{52}\}, \{c_2, c_5\} \cup \{s_{12}, s_{32}\}, \text{ and } \{c_3, c_6\} \cup \{s_{11}, s_{14}, s_{22}, s_{23}\}$$

consisting of the states on the cycle together with the states that are bisimilar, on the one hand, and the two blocks with the remaining states

$$\{s_{31}\} \cup \{s_{42}, s_{43}, s_{44}\} \text{ and } \{s_{41}\} \cup \{s_{51}, s_{53}\}$$

on the other hand.

Lemma 5.2. *Let $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ be a deterministic LTS.*

- (a) *If $|\mathcal{A}| = 1$ then $es(L)$ consists of all cycles in L .*
- (b) *Every $s \in S$ has a path to an end structure of L .*

Proof. (a) Since an end structure S' is closed under transitions, S' is a lasso. Because S' is minimal and non-empty, it follows that S' is a cycle.

(b) Let $U = \{t \in S \mid s \xrightarrow{w}^* t, w \in \mathcal{A}^*\}$ be the set of states reachable from state s . Then the set U is closed under transitions. The minimal non-empty subset $U' \subseteq U$ which is still closed under transitions is an end structure of L and can be reached by s . \square

Next we enhance the notion of a partition refinement algorithm. Now, an oracle can be consulted for the states in the end structures. In this approach, the initial partition is replaced by a partition in which all bisimilarity equivalence classes of states in end structures are split off from the original blocks.

Definition 5.3. A partition refinement algorithm with end structure oracle yields for an LTS $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ a valid refinement sequence $\Pi = (\pi'_0, \pi_1, \dots, \pi_n)$ where π'_0 is the end structure partition of L . The partition π'_0 is called the updated initial partition of L .

As Roberts' algorithm witnesses, in the case of a singleton action set the availability of an end structure oracle results in an algorithm with linear asymptotic performance. In the remainder of this section we confirm that in the case of more action labels the end structure does not help. The next lemma states that the amount of work required for the bisplitter \mathcal{B}_k by a partition refinement algorithm enhanced with an oracle, dealing with end structures, is at least the amount of work needed by a partition refinement algorithm without oracle for the bisplitter \mathcal{B}_{k-2} .

Lemma 5.4. *For the bisplitter $\mathcal{B}_k = (S, \mathcal{A}, \rightarrow, \pi_0)$, for some $k > 2$, let π'_0 be the updated initial partition. Then every valid refinement sequence $\Pi = (\pi'_0, \pi_2, \dots, \pi_n)$ for the updated bisplitter $\mathcal{B}'_k = (S, \mathcal{A}, \rightarrow, \pi'_0)$ satisfies $rc(\Pi) \geq rc(\mathcal{B}_{k-2})$.*

Proof. Observe that there are only two end structures in \mathcal{B}_k , viz. the singletons of the two states 0^k and 10^{k-1} . Since all other states can reach 0^k or 10^{k-1} , these states are not in an end structure: Choose $\sigma \in \mathbb{B}^k$, $\sigma \neq 0^k, 10^{k-1}$. Then σ is of the form $b0^j1\theta$ for some $b \in \mathbb{B}$, $j \geq 0$ and $\theta \in \mathbb{B}^*$. For $j = 0$ we have $\sigma \xrightarrow{a_1} b0^{k-1}$ which is either 0^k or 10^{k-1} ; for $j > 0$ we have $\sigma \xrightarrow{a_{j+1}} b0^{j-1}10^{k-(j+1)}$ while $b0^{j-1}10^{k-(j+1)}$ reaches 0^k or 10^{k-1} by induction.

By Lemma 4.3, every state $\sigma \in \mathbb{B}^k$ of \mathcal{B}_k has its own bisimulation equivalence class $\{\sigma\}$. It follows that the updated initial partition π'_0 consists of the blocks $\{0^k\}$, $\{10^{k-1}\}$, $B_0^k = B_0 \setminus \{0^k\}$, and $B_1^k = B_1 \setminus \{10^{k-1}\}$. Now, assume $\Pi = (\pi'_0, \pi_1, \dots, \pi_n)$ is a valid refinement sequence for \mathcal{B}_k . We show that $rc(\Pi) \geq rc(\mathcal{B}_{k-2})$ by constructing a valid refinement sequence Π' for \mathcal{B}_{k-2} satisfying $rc(\Pi) \geq rc(\Pi') \geq rc(\mathcal{B}_{k-2})$.

To construct Π' from Π , we use the partial projection function $p : \mathbb{B}^k \rightarrow \mathbb{B}^{k-2}$ that removes the prefix 11 from a bitstring and is undefined if 11 is not a prefix. That means $p(11\sigma) = \sigma$ for all $\sigma \in \mathbb{B}^{k-2}$ and $p(\sigma')$ is undefined for $\sigma' \notin 11\mathbb{B}^{k-2}$. A partition π of \mathbb{B}^k is projected to a partition of \mathbb{B}^{k-2} by projecting all the blocks of π and ignoring empty results, thus

$$p(\pi) = \{p[B] \mid B \in \pi\} \setminus \emptyset.$$

In particular, $p(\pi'_0) = \{\{\sigma \mid \sigma \in \mathbb{B}^{k-2}\}\}$, i.e. the unit partition of \mathbb{B}_{k-2} consisting of the prefix block B_ε only. Second, we remove repeated partitions from the sequence $(p(\pi'_0), p(\pi_1), \dots, p(\pi_n))$ to obtain a subsequence Π' , say $\Pi' = (\varrho_0, \varrho_1, \dots, \varrho_\ell)$. Thus, for

some order preserving surjection $q : \{1, \dots, n\} \rightarrow \{1, \dots, \ell\}$ it holds that $p(\pi_i) = p(\pi_{i'})$ iff $q(i) = q(i')$, and $\varrho_j = p(\pi_i)$ if $q(i) = j$ for $1 \leq i \leq n$, $1 \leq j \leq \ell$.

We have $\varrho_0 = p(\pi'_0) = \{B_\varepsilon\}$. Next we claim that $\varrho_1 = \pi_0^{k-2} = \{0\mathbb{B}^{k-3}, 1\mathbb{B}^{k-3}\}$ the initial partition of \mathcal{B}_{k-2} , containing the prefix blocks of 0 and 1 of \mathbb{B}^{k-2} : Suppose to the contrary that $b\theta, b\theta' \in \mathbb{B}^{k-2}$, for a bit $b \in \mathbb{B}$ and strings $\theta, \theta' \in \mathbb{B}^{k-3}$, are two different states which are not in the same block of ϱ_1 . Let i , $0 \leq i < n$ be such that $p(\pi_i) = \varrho_0$ and $p(\pi_{i+1}) = \varrho_1$. Then $11b\theta$ and $11b\theta'$ have been separated when refining π_i into π_{i+1} . But no action a_j witnesses such a split: (i) $\mathcal{B}_k(11b\theta, a_1) = \mathcal{B}_k(11b\theta', a_1)$ as both equal 0^k ; (ii) $\mathcal{B}_k(110\theta, a_2) = 110\theta \in B_1^k$, and $\mathcal{B}_k(110\theta', a_2) = 110\theta' \in B_1^k$; (iii) $\mathcal{B}_k(111\theta, a_2) = \mathcal{B}_k(111\theta', a_2) = 10^{k-1} \in B_1^k$; (iv) for $j > 2$ it holds that $\mathcal{B}_k(11b\theta, a_j), \mathcal{B}_k(11b\theta', a_j) \in B_1^k$. Since $\varrho_1 \neq \varrho_0$, ϱ_1 has at least two blocks. Hence, these must be $0\mathbb{B}^{k-3}$ and $1\mathbb{B}^{k-3}$. Thus $\varrho_1 = \{0\mathbb{B}^{k-3}, 1\mathbb{B}^{k-3}\}$ as claimed.

Next we prove that every refinement of ϱ_i into ϱ_{i+1} of Π' , for i , $1 \leq i < \ell$, is valid for \mathcal{B}_{k-2} . We first observe that, for all $\sigma, \sigma' \in \mathbb{B}^{k-2}$, $a_j \in \mathcal{A}$, it holds that $\mathcal{B}_{k-2}(\sigma, a_j) = \sigma'$ iff $\mathcal{B}_k(11\sigma, a_{j+2}) = 11\sigma'$. This is a direct consequence of the definition of the transition functions of \mathcal{B}_{k-2} and \mathcal{B}_k . From this we obtain

$$\sigma =_{\varrho_i} \sigma' \iff 11\sigma =_{\pi_h} 11\sigma' \quad (5.1)$$

provided $\varrho_i = p(\pi_h)$, for $0 \leq i \leq \ell$ and a suitable choice of h , via the definition of the projection function p . Now, consider the subsequent partitions ϱ_i and ϱ_{i+1} in Π' , $1 \leq i \leq \ell$. Now, let h , $0 \leq h < n$, be such that $\varrho_i = p(\pi_h)$ and $\varrho_{i+1} = p(\pi_{h+1})$. Clearly, ϱ_{i+1} is a refinement of ϱ_i ; if for $B \in \pi_{h+1}$ we have $B = \bigcup_{\alpha \in I} B_\alpha$ with $B_\alpha \in \pi_h$ for $\alpha \in I$, then for $p[B] \in \varrho_{i+1}$ we have $p[B] = \bigcup_{\alpha \in I} p[B_\alpha]$ with $p[B_\alpha] \in \varrho_i$ for $\alpha \in I$. The validity of the refinement of ϱ_i into ϱ_{i+1} is justified by the validity of π_h into π_{h+1} . If $\sigma =_{\varrho_i} \sigma'$ and $\sigma \neq_{\varrho_{i+1}} \sigma'$ for $\sigma, \sigma' \in \mathbb{B}^{k-2}$, then $\sigma, \sigma' \in 0\mathbb{B}^{k-3}$ or $\sigma, \sigma' \in 1\mathbb{B}^{k-3}$ since ϱ_i is a refinement of ϱ_0 . Moreover, $11\sigma =_{\pi_h} 11\sigma'$ and $11\sigma \neq_{\pi_{h+1}} 11\sigma'$ by (5.1). Hence, by validity, $\mathcal{B}_k(11\sigma, a_j) \neq_{\pi_h} \mathcal{B}_k(11\sigma', a_j)$ for some $a_j \in \mathcal{A}$. Clearly $j \neq 1$, since $(11\sigma)[2] = (11\sigma')[2] = 1$. Also, $j \neq 2$, since $\sigma[1] = \sigma'[1]$ we have $(11\sigma)[3] = (11\sigma')[3]$. Therefore, $\mathcal{B}_{k-2}(\sigma, a_{j-2}) \neq_{\varrho_i} \mathcal{B}_{k-2}(\sigma', a_{j-2})$, showing the refinement of ϱ_i into ϱ_{i+1} to be valid.

Finally, since every block in π_n is a singleton, this is also the case for ϱ_ℓ . Thus, ϱ_ℓ is indeed the coarsest stable partition for \mathcal{B}_{k-2} as required for Π' to be a valid refinement sequence for \mathcal{B}_{k-2} . Every refinement of ϱ_i into ϱ_{i+1} of Π' is projected from a refinement of some π_h into π_{h+1} of Π as argued above. Therefore, since $p(\pi_h) = \varrho_i$ and $p(\pi_{h+1}) = \varrho_{i+1}$, we have $rc(\pi_h, \pi_{h+1}) \geq rc(\varrho_i, \varrho_{i+1})$, and hence $rc(\Pi) = \sum_{h=1}^n rc(\pi_{h-1}, \pi_h) \geq \sum_{i=1}^\ell rc(\varrho_{i-1}, \varrho_i) = rc(\Pi')$. Since, by definition, $rc(\mathcal{B}_{k-2})$ is the minimum over all valid refinement sequences for \mathcal{B}_{k-2} it holds that $rc(\Pi') \geq rc(\mathcal{B}_{k-2})$. Therefore, $rc(\Pi) \geq rc(\Pi') \geq rc(\mathcal{B}_{k-2})$ as was to be shown. \square

Next we combine the above lemma with the lowerbound provided by Theorem 4.4 in order to prove the main result of this section.

Theorem 5.5. *Any partition refinement algorithm with an end structure oracle to decide bisimilarity for a deterministic LTS is $\Omega(n \log n)$.*

Proof. Let \mathcal{B}'_k be the updated bisplitter with the initial partition π'_0 containing $\{0^k\}$, $B_0 \setminus \{0^k\}$, $\{10^{k-1}\}$, and $B_1 \setminus \{10^{k-1}\}$ as given by the oracle for end structures rather than the partition π_0 containing B_0 and B_1 . By Lemma 5.4 we have, for $k > 2$, that $rc(\mathcal{B}'_k) \geq rc(\mathcal{B}_{k-2})$. By Theorem 4.4 we know that $rc(\mathcal{B}_{k-2}) \geq \frac{1}{2} \tilde{n} \log \frac{1}{2} \tilde{n}$ for $\tilde{n} = 2^{k-2}$, the number of states of \mathcal{B}_{k-2} . It holds that $\tilde{n} = \frac{2^{k-2}}{2^k} n = \frac{1}{4} n$. So $rc(\mathcal{B}'_k) \geq \frac{1}{8} n \log \frac{1}{8} n$ from which we

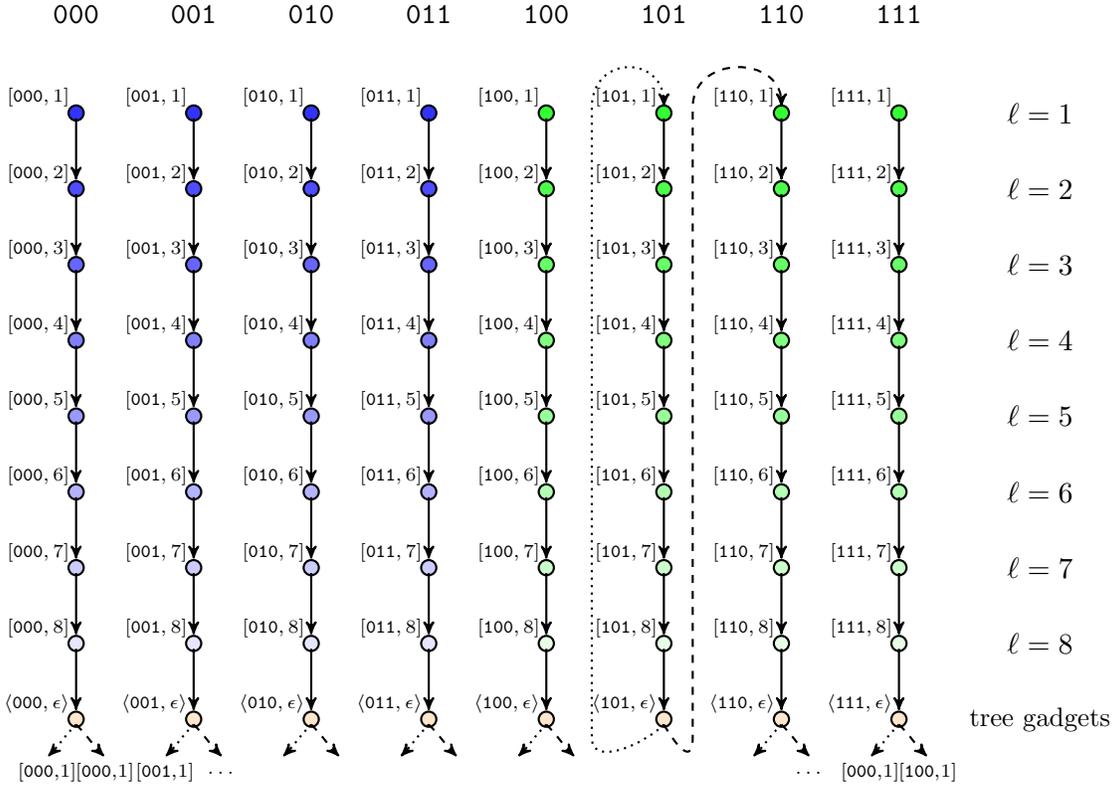


Figure 3: The partial layered bisplitter \mathcal{E}_3 with tree gadgets, the colours represent the initial partition.

conclude that deciding bisimilarity for \mathcal{B}_k with the help of an oracle for the end structures is $\Omega(n \log n)$. □

6. \mathcal{E}_k IS $\Omega((m + n) \log n)$ FOR PARTITION REFINEMENT

We modify the bisplitter \mathcal{B}_k , that has an action alphabet of $k-1$ actions, to obtain a deterministic LTS with two actions only. The resulting LTS \mathcal{E}_k has the action alphabet $\{a, b\}$, for each $k > 1$, and is referred to as the k -th *layered bisplitter*. We use \mathcal{E}_k to obtain an $\Omega((n + m) \log n)$ lowerbound for deciding bisimilarity for LTSs with only two actions, where n is the number of states and m is the number of transitions.

In order to establish the lowerbound we adapt the construction of \mathcal{B}_k at two places. We introduce for each $\sigma \in \mathbb{B}^k$, a stake of 2^k states. Moreover, to each stake we add a tree gadget. These gadgets have height $\lceil \log(\frac{k-1}{2}) \rceil$ to accommodate $\lceil (k-1)/2 \rceil$ leaves in order to encode the action alphabet \mathcal{A}_k of \mathcal{B}_k with $k-1$ actions.

Definition 6.1. Let $k > 1$, \mathcal{B}_k be the k -th bisplitter, and $\mathbb{A} = \{a, b\}$ be a two-element action set. The deterministic LTS $\mathcal{E}_k = (S_k^\mathcal{E}, \mathbb{A}, \rightarrow_\mathcal{E}, \pi_0^\mathcal{E})$, over the action set \mathbb{A} ,

- (a) has the set of states $S_k^\mathcal{E}$ defined as

$$S_k^{\mathcal{C}} = \{ [\sigma, \ell] \in \mathbb{B}^k \times \mathbb{N} \mid 1 \leq \ell \leq 2^k \} \cup \{ \langle \sigma, w \rangle \in \mathbb{B}^k \times \mathbb{A}^* \mid 0 \leq |w| \leq \lceil \log(\frac{k-1}{2}) \rceil \},$$

(b) has the transition relation $\rightarrow_{\mathcal{C}}$ given by

$$\begin{aligned} [\sigma, \ell] &\xrightarrow{\alpha}_{\mathcal{C}} [\sigma, \ell + 1] && \text{for } \sigma \in \mathbb{B}^k, 1 \leq \ell < 2^k, \alpha \in \mathbb{A} \\ [\sigma, 2^k] &\xrightarrow{\alpha}_{\mathcal{C}} \langle \sigma, \varepsilon \rangle && \text{for } \sigma \in \mathbb{B}^k, \alpha \in \mathbb{A} \\ \langle \sigma, w \rangle &\xrightarrow{\alpha}_{\mathcal{C}} \langle \sigma, w\alpha \rangle && \text{for } \sigma \in \mathbb{B}^k, |w| < \lceil \log(\frac{k-1}{2}) \rceil, \alpha \in \mathbb{A} \\ \langle \sigma, w \rangle &\xrightarrow{\alpha}_{\mathcal{C}} [\sigma', 1] && \text{for } \sigma \in \mathbb{B}^k, |w| = \lceil \log(\frac{k-1}{2}) \rceil, lbl(w\alpha) = j, \\ &&& \text{and, } \mathcal{B}_k(\sigma, a_j) = \sigma', \alpha \in \mathbb{A}, \end{aligned}$$

(c) and has the initial partition $\pi_0^{\mathcal{C}} = \{ C_0^{\ell}, C_1^{\ell} \mid 1 \leq \ell \leq 2^k \} \cup \{ C_{\varepsilon} \}$ defined as

$$\begin{aligned} C_0^{\ell} &= \{ [\sigma, \ell] \mid \sigma \in B_0 \} && \text{for } 1 \leq \ell \leq 2^k \\ C_1^{\ell} &= \{ [\sigma, \ell] \mid \sigma \in B_1 \} && \text{for } 1 \leq \ell \leq 2^k \\ C_{\varepsilon} &= \{ \langle \sigma, w \rangle \in S_k^{\mathcal{C}} \mid \sigma \in \mathbb{B}^k, w \in \mathbb{A}^* \}. \end{aligned}$$

The auxiliary labelling function $lbl : \mathbb{A}^{\leq \lceil \log(k-1) \rceil} \rightarrow \mathbb{N}$, used in item (b) is defined by $lbl(w) = \min\{bin(w)+1, k-1\}$. Here $bin : \mathbb{A}^* \rightarrow \mathbb{N}$ is the binary evaluation function defined by $bin(\varepsilon) = 0$, $bin(wa) = 2 * bin(w)$, and $bin(wb) = 2 * bin(w)+1$.

We see that with each string $\sigma \in \mathbb{B}^k$ we associate in \mathcal{C}_k as many as 2^k stake states $[\sigma, 1], \dots, [\sigma, 2^k]$, one for each level ℓ , $1 \leq \ell \leq 2^k$. The stake states are traversed from the top $[\sigma, 1]$ to bottom $[\sigma, 2^k]$ for each string σ over \mathbb{A} of length 2^k . The tree gadget, with states $\langle \sigma, w \rangle$ for bit sequences σ and strings w over \mathbb{A} , consists of a complete binary tree of height $\lceil \log(\frac{k-1}{2}) \rceil$ that hence has $\lceil (k-1)/2 \rceil$ leaves. Traversal down the tree takes a left child on action a from \mathbb{A} , a right child on action b from \mathbb{A} . Together with the two actions of \mathbb{A} , $k-1$ source-label pairs can be encoded, connecting the stake on top of the tree gadgets $k-1$ times with other stakes. To simulate a transition $\sigma \xrightarrow{a_j} \sigma'$ of \mathcal{B}_k in \mathcal{C}_k from a leaf of a tree gadget of σ to the top of the stake of σ' , we need to be at a leaf $\langle \sigma, w \rangle$ of the tree gadget of σ such that the combined string $w\alpha$ for $\alpha \in \mathbb{A}$ is the binary encoding according to lbl of the index j . An α -transition thus leads from the source $\langle \sigma, w \rangle$ to the target $[\sigma', 1]$ if $\sigma \xrightarrow{a_j} \sigma'$ in \mathcal{B}_k and $w\alpha$ corresponds to j . The partition

$$\pi_0^{\mathcal{C}} = \{ C_0^{\ell}, C_1^{\ell} \mid 1 \leq \ell \leq 2^k \} \cup \{ C_{\varepsilon} \}$$

distinguishes, for each level ℓ , the states at level ℓ of the stakes of strings starting with 0 in C_0^{ℓ} , the states of the stakes at level ℓ of strings starting with 1 in C_1^{ℓ} , and the states of the tree gadgets collected in C_{ε} .

Figure 3 depicts the layered 3-splitter \mathcal{C}_3 . Because also \mathcal{B}_3 has an action set of size 2 the tree gadgets only consist of the root node of the form $\langle \sigma, \varepsilon \rangle$. In Figure 2 we see that for bisplitter \mathcal{B}_3 we have $101 \xrightarrow{a_1} 101$ and $101 \xrightarrow{a_2} 110$. In Figure 3 we have transitions $\langle 101, \varepsilon \rangle \xrightarrow{a} [101, 1]$ and $\langle 101, \varepsilon \rangle \xrightarrow{b} [110, 1]$ (dotted and dashed, respectively). Colouring of nodes is used to represent the initial partition $\pi_0^{\mathcal{C}}$ that contains 17 blocks: for each level ℓ , $1 \leq \ell \leq 2^3$, $\pi_0^{\mathcal{C}}$ contains a block holding the four states of the stakes in C_0^{ℓ} on the left and a block with the four stake states in C_1^{ℓ} on the right, and lastly one block consisting of the eight tree states in C_{ε} at the bottom of the picture.

The 6-th bisplitter \mathcal{B}_6 has five actions, a_1 to a_5 . A tree gadget for the layered bisplitter \mathcal{C}_6 with corresponding outgoing transitions is drawn in Figure 4. The tree has height $\lceil \log((6-1)/2) \rceil = \lceil \log \frac{5}{2} \rceil = 2$, hence it has $2^2 = 4$ leaves. Since each leaf has two outgoing transitions,

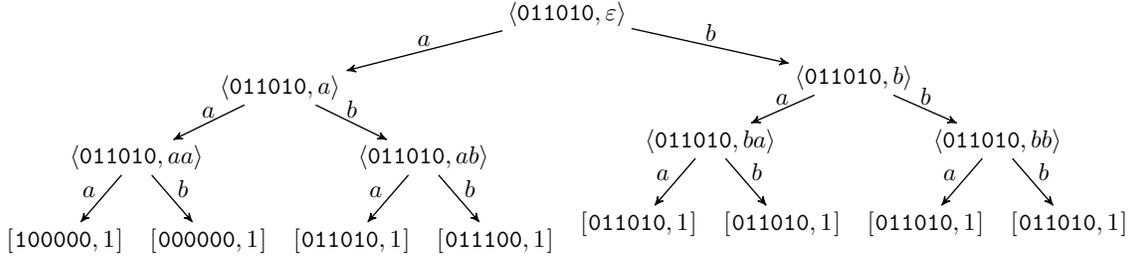


Figure 4: The example of the outgoing tree for \mathcal{C}_6 from the root $[011010, \varepsilon] \in S_6^{\mathcal{C}}$.

one labelled a and one labelled b , the two leftmost leaves $\langle \sigma, aa \rangle$ and $\langle \sigma, ab \rangle$ are used with the two labels a and b to simulate transitions for a_1 up to a_4 , the two rightmost leaves $\langle \sigma, ba \rangle$ and $\langle \sigma, bb \rangle$ have together four transitions all simulating the a_5 -transition of σ .

The next lemma introduces three facts for the layered bisplitter \mathcal{C}_k that we need in the sequel. The first states that if two states at different stakes, but at the same level, are separated during partition refinement, then all corresponding states at lower levels are separated as well. The second fact helps to transfer witnessing transitions in \mathcal{B}_k to the setting of \mathcal{C}_k . A transition $\sigma \xrightarrow{a_j} \sigma'$ of \mathcal{B}_k is reflected by a path from $[\sigma, 2^k]$ through the tree gadget of σ from root to leaf and then to the top state $[\sigma', 1]$ of the stake of σ' . The word $w\alpha$ encountered going down and out the tree gadget corresponds to the action a_j according to the lbl -function. Lastly, it is shown that no two pairs of different states within the stakes are bisimilar.

Lemma 6.2. *Let Π be a valid refinement sequence for \mathcal{C}_k and π a partition in Π .*

- (a) *If two states $[\sigma, \ell], [\sigma', \ell] \in S_k^{\mathcal{C}}$, for $1 \leq \ell \leq 2^k$, are in a different block of π , then all pairs $[\sigma, m], [\sigma', m] \in S$, for all levels $m, \ell \leq m \leq 2^k$, are in different blocks of π .*
- (b) *If $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ are split for π , then there are $w \in \mathbb{A}^*$, $\alpha \in \mathbb{A}$, and $\sigma'_1, \sigma'_2 \in \mathbb{B}^k$ such that*

$$[\sigma_1, 2^k] \xrightarrow{w}_{\mathcal{C}}^* \langle \sigma_1, w \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_1, 1] \quad \text{and} \quad [\sigma_2, 2^k] \xrightarrow{w}_{\mathcal{C}}^* \langle \sigma_2, w \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_2, 1]$$

with $[\sigma'_1, 1]$ and $[\sigma'_2, 1]$ in different blocks of π .

- (c) *If π is the last refinement in Π , it contains the singletons of $[\sigma, \ell]$ for $\sigma \in \mathbb{B}^k$ and $1 \leq \ell \leq 2^k$.*

Proof. (a) For a proof by contradiction, suppose the partition π is the first partition of Π that falsifies the statement of the lemma. So $\pi \neq \pi_0^{\mathcal{C}}$, since for the initial partition $\pi_0^{\mathcal{C}}$ the statement holds. Thus, π is a refinement of a partition π' in Π . So, there are two states $[\sigma, \ell], [\sigma', \ell] \in S_k^{\mathcal{C}}$ in different blocks of π while the states $[\sigma, \ell+1], [\sigma', \ell+1]$ are in the same block of π and hence of π' . Since $[\sigma, \ell]$ and $[\sigma', \ell]$ only have transitions to $[\sigma, \ell+1]$ and $[\sigma', \ell+1]$, respectively, that are in the same block π' , the refinement would not have been valid. We conclude that no falsifying partition π in Π exists and that the lemma holds.

(b) We first prove, that for all $w \in \mathbb{A}^*$, $|w| \leq \lceil \log(k-1) \rceil - 1$, if $\langle \sigma_1, w \rangle$ and $\langle \sigma_2, w \rangle$ are split in π , then there are $v \in \mathbb{A}^*$ and $\alpha \in \mathbb{A}$ such that $\langle \sigma_1, w \rangle \xrightarrow{v}_{\mathcal{C}}^* \langle \sigma_1, wv \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_1, 1]$ and $\langle \sigma_2, w \rangle \xrightarrow{v}_{\mathcal{C}}^* \langle \sigma_2, wv \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_2, 1]$ with $[\sigma'_1, 1]$ and $[\sigma'_2, 1]$ in different blocks of π . We prove this for all possible lengths $|w|$ by reverse induction. If w has maximal length, i.e. $|w| = \lceil \log(k-1) \rceil - 1$ this is clear. If $\langle \sigma_1, w \rangle$ and $\langle \sigma_2, w \rangle$ are split, for $|w| < \lceil \log(k-1) \rceil - 1$,

then either a -transitions or b -transitions lead to split states. By the induction hypothesis, suitable paths exist from the targets of such transitions. Adding the respective transition proves the induction hypothesis. Since $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ can only reach $\langle \sigma_1, \varepsilon \rangle$ and $\langle \sigma_2, \varepsilon \rangle$ the statement follows.

(c) Choose ℓ , $1 \leq \ell \leq 2^k$ and define the relation $R \subseteq S_k^{\mathcal{B}} \times S_k^{\mathcal{B}}$ such that $(\sigma_1, \sigma_2) \in R$ iff the stake states $[\sigma_1, \ell], [\sigma_2, \ell] \in S_k^{\mathcal{C}}$ are bisimilar for \mathcal{C}_k . We verify that R is a bisimulation relation for \mathcal{B}_k . Note, that R respects $\pi_k^{\mathcal{B}}$, the initial partition of \mathcal{B}_k . Now, suppose $(\sigma_1, \sigma_2) \in R$ and $\sigma_1 \xrightarrow{a_j} \sigma'_1$ for some $a_j \in \mathcal{A}_k$ and $\sigma'_1 \in S_k^{\mathcal{B}}$. By construction of \mathcal{C}_k we have

$$[\sigma_1, \ell] \xrightarrow{a_j}^* [\sigma_1, 2^k] \xrightarrow{a} \langle \sigma_1, \varepsilon \rangle \xrightarrow{w}^* \langle \sigma_1, w \rangle \xrightarrow{\alpha} [\sigma'_1, 1] \xrightarrow{a^{\ell-1}}^* [\sigma'_1, \ell]$$

where $lbl(w\alpha) = j$. Since $[\sigma_1, \ell]$ and $[\sigma_2, \ell]$ are bisimilar in \mathcal{C}_k , it follows that a corresponding path $[\sigma_2, \ell] \rightarrow^* [\sigma'_2, \ell]$ exists in \mathcal{C}_k with $[\sigma'_1, \ell]$ and $[\sigma'_2, \ell]$ bisimilar in \mathcal{C}_k . From this we derive that $\sigma_2 \xrightarrow{a_j} \sigma'_2$ in \mathcal{B}_k and $(\sigma'_1, \sigma'_2) \in R$. Hence, R is a bisimulation relation for \mathcal{B}_k indeed. It holds that bisimilarity of \mathcal{B}_k is the identity relation. Thus, if two stake states $[\sigma_1, \ell]$ and $[\sigma_2, \ell]$ are bisimilar for \mathcal{C}_k , then σ_1 and σ_2 are bisimilar for \mathcal{B}_k thus $\sigma_1 = \sigma_2$, and therefore $[\sigma_1, \ell] = [\sigma_2, \ell]$. \square

The next lemma states that the splitting of states $[\sigma, \ell] \in S^{\mathcal{C}}$, for each level ℓ , has refinement costs that are at least that of \mathcal{B}_k .

Lemma 6.3. *It holds that $rc(\mathcal{C}_k) \geq 2^k rc(\mathcal{B}_k)$ for all $k > 1$.*

Proof. Let $\Pi = (\pi_0^{\mathcal{C}}, \pi_1, \dots, \pi_n)$ be a valid refinement sequence for \mathcal{C}_k . We show that for each level ℓ , the sequence Π induces a valid refinement sequence Π^ℓ for \mathcal{B}_k .

For each $\ell \in \mathbb{N}$, such that $1 \leq \ell \leq m$, we define a partial projection function $p_\ell : S_k^{\mathcal{C}} \rightarrow \mathbb{B}^k$. The mapping p_ℓ maps states of shape $[\sigma, \ell] \in S_k^{\mathcal{C}}$ of \mathcal{C}_k to $\sigma \in \mathbb{B}^k$ and is undefined on all other states. A block B in a partition of \mathcal{C}_k is mapped to the block $p_\ell[B]$ of \mathbb{B}^k , by applying p_ℓ on all elements, resulting in:

$$p_\ell[B] = \{ \sigma \in \mathbb{B}^k \mid [\sigma, \ell] \in B \}.$$

A partition $p_\ell(\pi)$ of \mathcal{B}_k is obtained by applying p_ℓ to a partition π of \mathcal{C}_k and ignoring the empty blocks, i.e. $p_\ell(\pi) = \{ p_\ell[B] \mid B \in \pi \} \setminus \emptyset$. The sequence $\Pi^\ell = (\pi_0^\ell, \dots, \pi_m^\ell)$ is obtained from the sequence $(p_\ell(\pi_0^{\mathcal{C}}), p_\ell(\pi_1), \dots, p_\ell(\pi_n))$ by removing possible duplicates. We verify that Π^ℓ is a valid refinement sequence for \mathcal{B}_k .

First, we check that π_i^ℓ is a refinement of π_{i-1}^ℓ , for $1 \leq i \leq m$. Choose index i arbitrarily. Let the index h with $1 \leq h \leq n$ be such that $p_\ell(\pi_{h-1}) = \pi_{i-1}^\ell$ and $p_\ell(\pi_h) = \pi_i^\ell$. Then we fix a block $B \in \pi_i^\ell$. Since $\pi_i^\ell = p_\ell(\pi_h)$ there is a block $B' \in \pi_h$ such that $B = p_\ell[B']$. Since π_h is a refinement of π_{h-1} there is a block $B'' \in \pi_{h-1}$ such that $B' \subseteq B''$. This implies that $p_\ell[B'] \subseteq p_\ell[B'']$ and since $B = p_\ell[B'] \neq \emptyset$ also $p_\ell[B''] \neq \emptyset$. So, we conclude that $p_\ell[B''] \in \pi_{i-1}^\ell$ and $B \subseteq p_\ell[B'']$. Thus π_i^ℓ is a refinement of π_{i-1}^ℓ .

Next, we verify that Π^ℓ is a valid refinement sequence for \mathcal{B}_k . Suppose the state $\sigma_1, \sigma_2 \in S_k^{\mathcal{B}}$ are split for the refinement of π_{i-1}^ℓ into π_i^ℓ . Then the states $[\sigma_1, \ell], [\sigma_2, \ell] \in S_k^{\mathcal{C}}$ are split for the refinement of a partition π_{h-1} into the partition π_h for some index h , with $1 \leq h \leq n$. Then either (i) $\ell = 2^k$ and $[\sigma_1, \ell]$ and $[\sigma_2, \ell]$ have α -transitions to different blocks, for some $\alpha \in \mathbb{A}$, or (ii) $\ell < 2^k$ and $[\sigma_1, \ell+1]$ and $[\sigma_2, \ell+1]$ are in different blocks of π_{h-1} . In the case of (ii), it follows by Lemma 6.2 that also $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ are in different blocks of π_{h-1} . Thus, the refinement of some π_{g-1} into π_g , $1 \leq g \leq h \leq n$, split the two states

$[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$. By Lemma 6.2 there are $w \in \mathbb{A}^*$, $\alpha \in \mathbb{A}$, and $\sigma'_1, \sigma'_2 \in \mathbb{B}^k$ such that

$$[\sigma_1, 2^k] \xrightarrow{w}_{\mathcal{C}}^* \langle \sigma_1, w \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_1, 1] \quad \text{and} \quad [\sigma_2, 2^k] \xrightarrow{w}_{\mathcal{C}}^* \langle \sigma_2, w \rangle \xrightarrow{\alpha}_{\mathcal{C}} [\sigma'_2, 1]$$

with $[\sigma'_1, 1]$ and $[\sigma'_2, 1]$ in different blocks of π_{g-1} . Hence, σ'_1 and σ'_2 are in different blocks of π_{i-1}^ℓ while $\sigma_1 \xrightarrow{a_j}_{\mathcal{B}} \sigma'_1$ and $\sigma_2 \xrightarrow{a_j}_{\mathcal{B}} \sigma'_2$ for $j = \text{lbl}(w\alpha)$, which justifies splitting σ_1 and σ_2 for π_j^ℓ . We conclude that Π^ℓ is a valid refinement sequence for \mathcal{B}_k .

We have established that if Π is a valid refinement sequence for \mathcal{C}_k , then Π^ℓ is a valid refinement sequence for \mathcal{B}_k . The sequence Π^ℓ is obtained from Π by sifting out the blocks of Π 's partitions and removing repeated partitions. Therefore it holds that $rc(\Pi) \geq rc(\Pi^\ell)$. Since the mappings p_ℓ and $p_{\ell'}$ include pairwise distinct sets of stake states for $\ell \neq \ell'$, $1 \leq \ell \leq 2^k$, it follows that $rc(\Pi) \geq \sum_{\ell=1}^{2^k} rc(\Pi^\ell) \geq 2^k rc(\mathcal{B}_k)$. Taking the minimum over all valid refinement sequences for \mathcal{C}_k we conclude that $rc(\mathcal{C}_k) \geq 2^k rc(\mathcal{B}_k)$ as was to be shown. \square

With the above technical lemma in place, we are able to strengthen the $\Omega(n \log n)$ lowerbound of Theorem 4.4 by now taking the number of transitions into account. The improved lowerbound is $\Omega((m+n) \log n)$, where m is the number of transitions and n the number of states.

Theorem 6.4. *Deciding bisimilarity for (deterministic) LTSs with a partition refinement algorithm is $\Omega((m+n) \log n)$, where n is the number of states and m is the number of transitions of the LTS.*

Proof. For the bisplitter \mathcal{B}_k , we know by Theorem 4.4 that $rc(\mathcal{B}_k) \geq 2^{k-1}(k-1)$. Thus, by Lemma 6.3, we obtain $rc(\mathcal{C}_k) \geq 2^{2k-1}(k-1)$. In the case of \mathcal{C}_k we have for n and m that $n = 2^k(2^k + 2^{\lceil \log(k-1) \rceil} - 1)$ and $m = 2n$. Hence $n + m \in \Theta(2^{2k-1})$ and $\log n \in \Theta(k-1)$, from which it follows that $rc(\mathcal{C}_k) \in \Omega((m+n) \log n)$. \square

Underlying the proof of the lowerbound for deciding bisimilarity for the family of layered bisplitters \mathcal{C}_k is the observation that each \mathcal{C}_k can be seen as 2^k stacked instances of the ordinary bisplitters \mathcal{B}_k , augmented with tree gadgets to handle transitions properly. The other essential ingredient for the proof of Theorem 6.4 is the complexity of deciding bisimilarity with a partition refinement algorithm on the \mathcal{B}_k family. The same reasoning applies when considering partition refinement algorithms with an oracle for end structures from Section 5. Also with an oracle the lowerbound of $\Omega((m+n) \log n)$ remains.

Theorem 6.5. *Any partition refinement algorithm with an oracle for end structures that decides bisimilarity for (deterministic) LTSs is $\Omega((m+n) \log n)$.*

Proof. The proof is similar to that of Lemma 5.4 and Theorem 6.4. Consider, for some $k > 2$, the layered bisplitter \mathcal{C}_k having initial partition $\pi_0^\mathcal{C}$. The LTS \mathcal{C}_k has two end structures, viz. the set $S_0 \subseteq S_k^\mathcal{C}$ containing the states of the stake and accompanying tree gadget $S_0 = \{ [0^k, \ell] \mid 1 \leq \ell \leq 2^k \} \cup \{ \langle 0^k, w \rangle \mid w \in \mathbb{A}^*, |w| \leq \lceil \log(\frac{k-1}{2}) \rceil \}$ for 0^k and a similar end structure $S_1 \subseteq S_k^\mathcal{C}$ for 10^{k-1} . The sets S_0 and S_1 are minimally closed under the transitions of \mathcal{C}_k . Other states, on the stake or tree gadget for a string σ , have a path to these sets inherited from a path from σ to 0^k or 10^k in \mathcal{B}_k . The bisimulation classes S'_0 and S'_1 , say, with respect to $S_k^\mathcal{C}$ rather than $\pi_0^\mathcal{C}$, consist of S_0 and S_1 themselves plus a part of the tree gadgets for transitions in \mathcal{C}_k leading to S_0 and S_1 , respectively.

The update of the initial partition $\pi_0^\mathcal{C}$ with oracle information, which concerns, ignoring the tree gadgets, the common refinement of the layers $\{ [\sigma, \ell] \mid \sigma \in B_0 \}$ and $\{ [\sigma, \ell] \mid \sigma \in B_1 \}$

on the one hand, and the bisimulation classes S'_0 and S'_1 on the other hand, is therefore equal to $\pi_0^{\mathcal{C}}$ on the stakes, and generally finer on the tree gadgets.

Next, every valid refinement sequence $\Pi = (\pi'_0, \pi_2, \dots, \pi_n)$ for the updated LTS $\mathcal{C}'_k = (S, \mathcal{A}, \rightarrow, \pi'_0)$ satisfies $rc(\Pi) \geq rc(\mathcal{C}_{k-2})$. Following the lines of the proof of Lemma 5.4, we can show that a valid refinement sequence Π for \mathcal{C}_k with updated initial partition π'_0 induces a valid refinement sequence Π' for \mathcal{C}_{k-2} .

The number of states in \mathcal{C}_{k-2} is $\Theta(n)$ with n the number of states of \mathcal{C}_k , and the number of transitions in \mathcal{C}_{k-2} is $\Theta(m)$ with m the number of transitions of \mathcal{C}_k . Therefore, $rc(\Pi) \geq rc(\Pi') \geq rc(\mathcal{C}_{k-2})$, from which we derive that any partition refinement algorithm with an oracle for end structures involves $\Theta((m+n) \log n)$ times moving a state for \mathcal{C}_k , and hence, the algorithm is $\Omega((m+n) \log n)$. \square

7. AN $\Omega(n)$ LOWERBOUND FOR PARALLEL PARTITIONING ALGORITHMS

In this section we pose the question of the effect of the concept of valid refinements on parallel partition refinement algorithms. We show an $\Omega(n)$ lowerbound. This result was already suggested in [Kul13, Theorem 3], without making explicit which operations are allowed to calculate the refinement. In particular, for deterministic LTSs with singleton alphabets, an $O(\log n)$ parallel refinement algorithm [JR94] exists, defying the argumentation of [Kul13]. This latter algorithm clearly is not based on valid refinements.

Parallel bisimulation algorithms are most conveniently studied in the context of PRAMs (Parallel Random Access Machine) [SV84], which have an unbounded number of processors that can all access the available memory. PRAMs are approximated by GPUs (Graphical Processing Units) that currently contain thousands of processor cores, but more interestingly, in combination with the operating system, can run millions of independent threads simultaneously.

There are a few variants of the PRAM model. The most important variation is in what happens when multiple processors try to write to the same address in memory. In the *common* scheme a write to a particular address takes place if all processors writing to this address write the same value. Otherwise, the write fails and the address will contain an arbitrary value. In the *arbitrary* scheme, one of the processors writing to the address will win, and writes its value; the writes of other processors to the address are ignored. In the *priority* scheme, the processor with the lowest index writes to the address.

A number of algorithms have been proposed to calculate bisimulation on PRAMs or GPUs [LR94, RL98, Wij15, MGH⁺21], and there are also parallel algorithms developed for networks of parallel computers [BO05]. The algorithms in [LR94, RL98] require $O(n \log n)$ time on respectively $\frac{m}{\log n} \log \log n$ and $\frac{m}{n} \log n$ processors. The algorithm in [MGH⁺21] has the best worst-case time complexity of $O(n)$ and uses $\max(n, m)$ processors. All these parallel algorithms have in common that they can be classified as partition refinement algorithms in the sense that they all calculate a valid sequence of partitions.

Note that parallel refinement algorithms can fundamentally outperform sequential algorithms. In order to understand why parallel algorithms achieve an upperbound of $O(n)$ vs. a lowerbound of $\Omega((m+n) \log n)$ for sequential algorithms, we look at the algorithm in [MGH⁺21] in more detail as it has the best time complexity. For the sake of exposition we assume here that there is only one action, although the story with multiple actions is essentially the same. In the algorithm, first an unstable block is chosen. All states reaching

this block are marked, which is done in constant time, by one processor per transition. Subsequently, all marked states in each block separate themselves from the other states in constant time using one processor per state, by employing an intricate trick where each block is characterised by a unique ‘leader’ state in each block. Here it is essential that the PRAM model uses the arbitrary or priority writing scheme. The algorithm does not work in the common scheme. In [MGH⁺21] it is shown that at most $3n$ of these constant-time splitting steps need to be performed, leading to a time complexity of $O(n)$.

An important observation is that parallel algorithms allow to split blocks in constant time, whereas for sequential algorithms we defined the refinement costs as the minimal number of states that had to be moved from one block to a new block. We assume, similarly to the sequential setting, that a new refinement can only be calculated on the basis of a previous refinement. This makes it natural to define the notion of parallel refinement costs as the minimal conceivable length of a valid refinement sequence and take this as the minimal time required to calculate bisimulation using partitioning in a parallel setting.

For an LTS L and a sequence $\Pi = (\pi_0, \dots, \pi_n)$, the *parallel refinement cost* is the number of refinements in the sequence, $prc(\Pi) = n$. For an LTS L we define

$$prc(L) = \min\{ prc(\Pi) \mid \Pi \text{ is a valid refinement sequence for } L \}.$$

Observe that parallel refinement costs allow for extremely fast partitioning of transition systems. Below we show an example with $2^k + k$ states with a refinement cost of 1. The states are given by $b_0, \dots, b_{k-1}, a_0, \dots, a_{2^k-1}$. There is a transition from a_i to b_j if the j -th bit in the binary representation of i is 1. The initial partition π_{init} groups all states a_i in one partition, and puts each state b_j in a partition of its own. So, π_{init} contains $k+1$ blocks. In Figure 5 this transition system is depicted for $k = 3$.

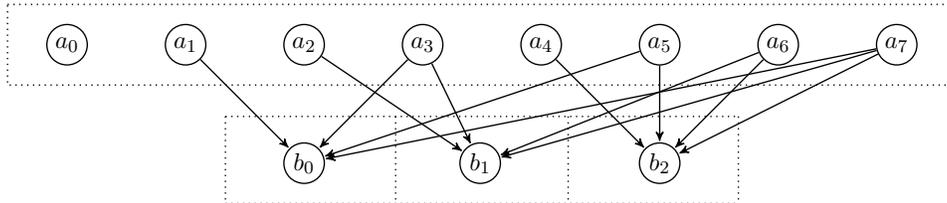


Figure 5: A transition system with a parallel refinement cost of 1

The shortest valid refinement sequence is $(\pi_{init}, \pi_{final})$, where in π_{final} each state is in a separate block. This refinement is valid, because in π_{init} there is enough information to separate each state from any other, as can easily be checked against the definition. As this refinement sequence has length 1, the parallel refinement cost of this transition system is 1, indicating that it is conceivable to make a bisimulation partitioning algorithm doing this refinement in constant time. Note that existing parallel algorithms do not achieve this performance. For instance, the algorithm in [MGH⁺21] requires linear time as it checks stability for each new block sequentially.

Although parallel partitioning can be fast, we show, using the notion of parallel refinement costs, that calculating bisimulation in parallel requires time $\Omega(n)$. For this purpose, we construct a family of LTSs \mathcal{D}_n for which the length of any valid refinement sequence grows linearly with the number of states.

Definition 7.1. For $n > 2$, the sequential splitter $\mathcal{D}_n = (S, \{a\}, \rightarrow, \pi_0)$ is defined as the LTS that has the set $S = \{1, \dots, n\}$ as its set of states, the relation

$$\rightarrow = \{(i, i+1) \mid 1 \leq i < n\} \cup \{(n, n)\}$$

as the transition relation, and the set $\pi_0 = \{\{1, \dots, n-1\}, \{n\}\}$ as its initial partition.

For every $n > 2$, the deterministic LTS $\mathcal{D}_n = (S, \Sigma, \rightarrow, \pi_0)$ has n states and n transitions. For $n = 8$ the transition system is depicted in Figure 6. The set of states $S = \{1, \dots, n\}$ form a chain in which every state $i \in \{1, \dots, n-1\}$ has an outgoing transition to the next state $(i, i+1) \in \rightarrow$. The state at the end of the chain $n \in S$ has a self loop $(n, n) \in \rightarrow$. In the initial partition there are two blocks, one block containing $n \in S$ and the other block containing all other states.

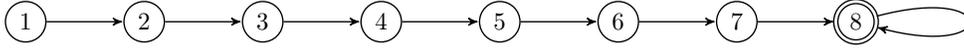


Figure 6: Sequential splitter \mathcal{D}_8 with initial partition $\{\{1, 2, 3, 4, 5, 6, 7\}, \{8\}\}$

The following lemma states that for every sequential splitter \mathcal{D}_n there is a unique valid refinement sequence.

Lemma 7.2. *For every $n > 2$, \mathcal{D}_n has a unique valid refinement sequence Π_n that consists of $n-1$ partitions.*

Proof. For $n > 2$ the sequential splitter \mathcal{D}_n has a valid partition refinement sequence which is given by $\Pi_{\mathcal{D}_n} = (\pi_1, \dots, \pi_{n-1})$, where $\pi_i = \{\{1, \dots, n-i\}, \{n-i+1\}, \{n-i+2\}, \dots, \{n\}\}$. This is proven by induction on the index i of the partition π_i .

Next, we must show that this refinement sequence is unique. So, in order to obtain a contradiction, consider some π_i and assume some valid refinement π'_{i+1} different from π_{i+1} exists. This means that there must either be two states $j, j' \in \{1, \dots, n-i-1\}$ that are in different blocks in π'_{i+1} , or state $n-i$ must be in the same block as some state $j < n-i$. In the first case states j and j' are in the same block in π_i and have exactly the same transitions to the same block in π_i . Hence, the states j and j' are in the same block in π_{i+1} . In the second case, all states $j' < n-i$ must be in the same block as j , using exactly the argument of the first case. The state $n-i$ is also in that same block. But then π'_{i+1} is not a strict refinement of π_i , making it invalid.

Thus, for each i , $1 \leq i \leq n-1$, no other valid refinement than π_{i+1} of π_i exists, making Π_n the only valid refinement sequence for \mathcal{D}_n . \square

The observation in Lemma 7.2 leads to the following theorem on the time complexity for parallel partition refinement algorithms.

Theorem 7.3. *Any parallel partition refinement algorithm that decides bisimilarity for an LTS with n states has time complexity $\Omega(n)$.*

Proof. For every $n \geq 2$ the LTS \mathcal{D}_n has n states and n transitions. Any algorithm that is a parallel partition refinement algorithm has time complexity that is at least the length of the shortest refinement sequence. According to Lemma 7.2 \mathcal{D}_n has a unique refinement sequence which witnesses $\text{prc}(\mathcal{D}_n) = n-1 \in \Omega(n)$. \square

The theorem puts a bound on the fastest possible parallel, partition based algorithms for bisimulation. But it must be observed that other techniques than partition refinement can produce faster algorithms, although it may be in more restricted settings. Concretely, the algorithm [JR94] that is based on the sequential Roberts' algorithm from Section 3, runs in $O(\log n)$ parallel time for deterministic transition systems with n states and only one transition label.

In order to obtain an idea of how these fast parallel algorithms work, we illustrate one of the major techniques to determine a bisimulation refinement of \mathcal{D}_n in time $O(\log n)$. In \mathcal{D}_n the initial partition has a small block $B_{small} = \{n\}$ and a large block $B_{large} = \{1, \dots, n-1\}$. States in B_{large} with a different number of steps to B_{small} cannot be bisimilar and can be split. Note that splitting on the basis of this distance is not a valid refinement in the sense of Definition 2.2.

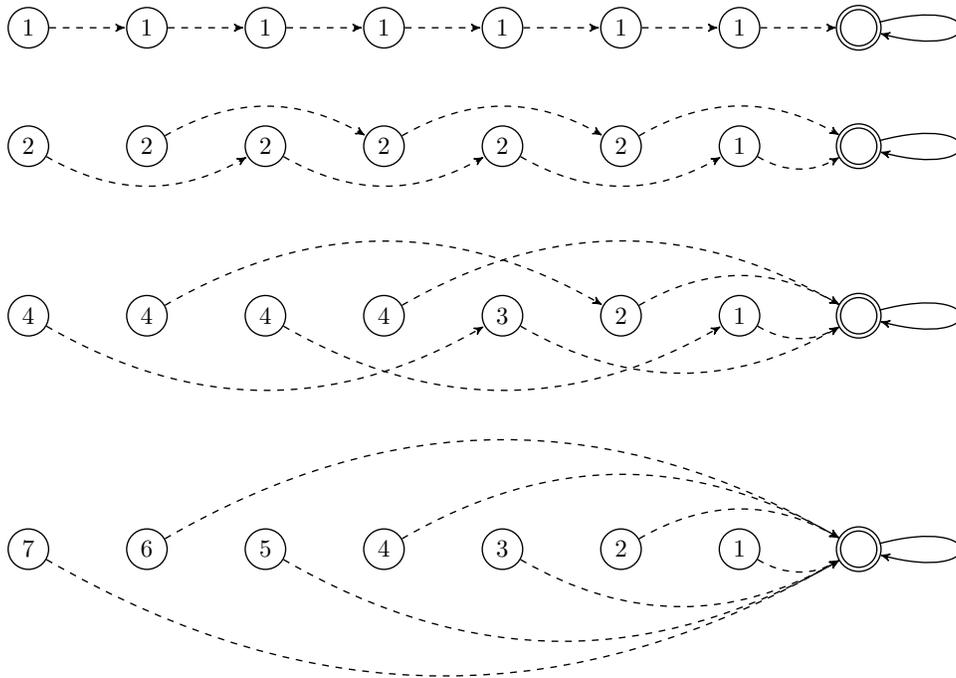


Figure 7: Calculating the distance to the rightmost state in $O(\log n)$ time

Determining the distance of the states in B_{large} to B_{small} can be done in $O(\log n)$ time in parallel, cf. [HJ86]. The basic idea is explained in Figure 7. Each state in the block at the left gets weight 1. We desire to sum up in each state the weights of all states to its right. We do this by adding up the weight of the right neighbour and adapting the outgoing transition to point to the state to which the right neighbour is pointing. We use dashed transitions to stress that we are now using the transitions for another purpose. In each round k a state contains the sum of all 2^k states to its right. So, after $\log_2 k$ rounds it contains the sum of all k states to its right, which is the distance to B_{small} .

An interesting open question is whether the notion of a valid refinement sequence can be adapted, such that the $\Omega(n)$ lowerbound would still apply when techniques such as parallel counting formulated above would be incorporated in the bisimulation algorithm.

8. COLOUR REFINEMENT

As for establishing bisimilarity on labelled transition systems, partition refinement is the standard approach for algorithms that do colour refinement. Colour refinement, also known as naive vertex classification or 1-dimensional Weisfeiler-Lehman test, is frequently applied, among others in the setting of deciding graph isomorphism [GN21]. Given a graph where each node has been assigned an initial colour, colour refinement asks to find a refining colouring with the least number of colours possible such that two nodes of the same colour have, for all colours, the same number of neighbours of the latter colour. Also for colour refinement, algorithms typically search for the coarsest stable colouring. See [GKMS21] for an overview.

Colour refinement is known to be $\Omega((m+n) \log n)$ with n the number of nodes and m the number of edges. Early algorithms of complexity $O((m+n) \log n)$ include [CC82, PT87]. The lowerbound for partition refinement algorithms has been established in [Ber15, BBG17]. However, for colour refinement the costs are measured in terms of inspected edges and hence are different from the costs for computing bisimilarity by partition refinement. The cost function underlying the complexity of colour refinement is defined by

$$\text{cost}(R, S) = |\{ (u, v) \in E \mid u \in R, v \in S \}|,$$

i.e. the number of transitions between the blocks R and S (sets including all nodes of chosen colours), where block R is recoloured in view of the colours in block S .

The paper [BBG17] provides a detailed implementation of an efficient algorithm with time complexity $O((m+n) \log n)$, that given a graph $G = (V, E)$ and colouring α finds the coarsest stable colouring refining α and compares the algorithm to other proposals in the literature. In this setting, a colouring $\gamma : V \rightarrow \mathbb{N}$ of G is stable iff $|\mathcal{N}(u) \cap \gamma^{-1}(c)| = |\mathcal{N}(v) \cap \gamma^{-1}(c)|$ for all nodes $u, v \in V$ and all colours $c \in \mathbb{N}$. Here, $\mathcal{N}(u)$ and $\mathcal{N}(v)$ denote the sets of neighbours of u and v in G , respectively. In order to establish a lowerbound, Berkholz et al. define a family of graphs \mathcal{G}_k for which the costs of computing the coarsest stable colouring is $\Omega((m+n) \log n)$ starting from the unit colouring assigning to all nodes the same colour. The paper also discusses the connection of colour refinement with equivalence in 2-variable logic and with finding bisimilarity on Kripke structures. Regarding the latter the focus is on Kripke structures rather than labelled transition systems, as in the present paper. In order to transfer the lowerbound result for colour refinement to a lowerbound result for bisimilarity by partition refinement for Kripke structures, the stability requirement for colouring mentioned above is adapted, viz. to $\mathcal{N}^+(u) \cap S = \emptyset \iff \mathcal{N}^+(v) \cap S = \emptyset$ for all blocks R, S and nodes u, v in R , with $\mathcal{N}^+(u)$ and $\mathcal{N}^+(v)$, i.e. the directly reachable states for u and v , respectively.

Central to the resulting family of Kripke structures \mathcal{S}_k in [BBG17] are complete bipartite graphs $K_{k,k}$ with k^2 transitions, one for each bitstring in \mathbb{B}^k , which are dense with respect to transitions. This is because the refinement costs incurred for colour refinement are based on counting edges. The family of LTSs \mathcal{C}_k presented in Section 6 has as their main components thin stakes that are 2^k states high, one for every bitstring in \mathbb{B}^k because for partition refinement for bisimulation the number of (moved) states is relevant.

Although the families of graphs for the lowerbounds of bisimulation and graph colouring are definitely related, they are very different if it comes to the approach, in particular regarding the measurement of complexity. It is unclear how to transform the \mathcal{S}_k -family into a family of labelled transition systems such that partition refinement takes $\Omega((m+n) \log n)$ transfers of states to a newly created block, i.e. in terms of the refinement cost rc of Section 2.

Still the Kripke structures \mathcal{S}_k of [BBG17] can be interpreted as non-deterministic labelled transition systems with a single action label and an initial partition based on the assignment of atomic proposition. Similarly, it is not obvious how to transform the \mathcal{C}_k -family into a family of undirected or directed graphs such that colour refinement requires inspection of $\Omega((m+n) \log n)$ edges.

9. CONCLUSION

We have shown that, even when restricted to deterministic LTSs, it is not possible to construct a sequential algorithm based on partition refinement that is more efficient than $\Omega((m+n) \log n)$. The bound obtained is preserved even when the algorithm is extended with an oracle that can determine for specific states in constant time whether they are bisimilar or not. The oracle proof technique enabled us to show that the algorithmic ideas underlying Roberts' algorithm [PTB85] for the one-letter alphabet case cannot be used to come up with a fundamentally faster enhanced partition refinement algorithm for bisimulation.

Of course, this is not addressing a generic lower bound to decide bisimilarity on LTSs, nor proving the conjecture that the Paige-Tarjan algorithm is optimal for deciding bisimilarity. It is conceivable that a more efficient algorithm for bisimilarity exists that is not based on partitioning. However, as it stands, no techniques are known to prove such a generic algorithmic lowerbound, and all techniques that do exist make assumptions on allowed operations, such as the well-known lowerbound on sorting.

But by relaxing the notion of a valid partition sequence, and maybe introducing alternatives for oracles, it may very well be possible that the lower bound is extended to a wider range of algorithmic techniques to determine bisimulation, making it very unlikely that sequential algorithms for bisimulation with a time-complexity better than $O((m+n) \log n)$ exist. Note that the current lowerbound already applies to all known efficient algorithms for bisimulation.

For the parallel setting, we showed that deciding bisimilarity by partitioning is $\Omega(n)$. In this case a similar situation occurs. For LTSs with one action label it is possible to calculate bisimulation in logarithmic time, cf. [JR94]. An interesting, but as yet open question is whether the techniques used [JR94] can fundamentally improve the efficiency of determining bisimulation in parallel, or, as we believe, the lowerbound result can be strengthened along the lines of the sequential case to show that the techniques of [JR94] are insufficient to obtain a sub-linear parallel time complexity to determine bisimulation for labelled transition systems with at least two action labels.

REFERENCES

- [BBG17] C. Berkholtz, P. Bonsma, and M. Grohe. Tight lower and upper bounds for the complexity of canonical colour refinement. *Theory of Computing Systems*, 60(4):581–614, 2017. doi:10.1007/s00224-016-9686-0.
- [BC04] J. Berstel and O. Carton. On the complexity of Hopcroft's state minimization algorithm. In M. Domaratzki et al., editor, *Proc. CIAA 2004*, volume 3317 of *Lecture Notes in Computer Science*, pages 35–44. Springer, 2004. doi:10.1007/978-3-540-30500-2_4.
- [Ber15] C. Berkholtz. *Lower Bounds for Heuristic Algorithms*. PhD thesis, RWTH Aachen, 2015.
- [BO05] S. Blom and S. Orzan. A distributed algorithm for strong bisimulation reduction of state spaces. *Software Technology for Technology Transfer*, 7(1):74–86, 2005. doi:10.1007/s10009-004-0159-4.

- [Buc99] P. Buchholz. Exact performance equivalence: An equivalence relation for stochastic automata. *Theoretical Computer Science*, 215:263–287, 1999. doi:10.1016/S0304-3975(98)00169-8.
- [CC82] A. Cardon and M. Crochemore. Partitioning a graph in $O(|A|\log_2|V|)$. *Theoretical Computer Science*, 19(1):85–98, 1982. doi:10.1016/0304-3975(82)90016-0.
- [CRS08] G. Castiglione, A. Restivo, and M. Sciortino. Hopcroft’s algorithm and cyclic automata. In C. Martín-Vide et al., editor, *Proc. LATA 2008*, volume 5196 of *Lecture Notes in Computer Science*, pages 172–183. Springer, 2008. doi:10.1007/978-3-540-88282-4_17.
- [DPP04] A. Dovier, C. Piazza, and A. Policriti. An efficient algorithm for computing bisimulation equivalence. *Theoretical Computer Science*, 311:221–256, 2004. doi:10.1016/S0304-3975(03)00361-X.
- [GKMS21] M. Grohe, K. Kersting, M. Mladenov, and P. Schweitzer. Color refinement and its applications. In G. Van den Broek, K. Kersting, and Natarajan S., editors, *An Introduction to Lifted Probabilistic Inference*, chapter 15. The MIT Press, 2021. doi:10.7551/mitpress/10548.003.0023.
- [GMV21] J.F. Groote, J. Martens, and E.P. de Vink. Bisimulation by partitioning is $\Omega((m+n)\log n)$. In S. Haddad and D. Varacca, editors, *Proc. CONCUR 2021*, volume 203 of *LIPICs*, pages 31:1–31:16. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2021. doi:10.4230/LIPICs.CONCUR.2021.31.
- [GN21] M. Grohe and D. Neuen. Isomorphism, canonization, and definability for graphs of bounded rank width. *Communications of the ACM*, 64(5):98–105, 2021. doi:10.1145/3453943.
- [GVV18] J.F. Groote, H.J. Rivera Verduzco, and E.P. de Vink. An efficient algorithm to determine probabilistic bisimulation. *Algorithms*, 11(9):131,1–22, 2018. doi:10.3390/a11090131.
- [HJ86] W.D. Hillis and G.L. Steele Jr. Data parallel algorithms. *Communications of the ACM*, 29(12):1170–1183, 1986. doi:10.1145/7902.7903.
- [Hop71] J. Hopcroft. An $n\log n$ algorithm for minimizing states in a finite automaton. In Z. Kohavi and A. Paz, editors, *Theory of Machines and Computations*, pages 189–196. Academic Press, 1971. doi:10.1016/b978-0-12-417750-5.50022-1.
- [JGKW20] D.N. Jansen, J.F. Groote, J.J.A. Keiren, and A. Wijs. An $O(m\log n)$ algorithm for branching bisimilarity on labelled transition systems. In A. Biere and D. Parker, editors, *Proc. TACAS*, volume 12079 of *Lecture Notes in Computer Science*, pages 3–20. Springer, 2020. doi:10.1007/978-3-030-45237-7_1.
- [JR94] J. Jájá and Kwan Woo Ryu. An efficient parallel algorithm for the single function coarsest partition problem. *Theoretical Computer Science*, 129(2):293–307, 1994. doi:10.1016/0304-3975(94)90030-2.
- [KMP77] D.E. Knuth, J.H. Morris, and V.R. Pratt. Fast pattern matching in strings. *SIAM Journal on Computing*, 6(2):323–350, 1977. doi:10.1137/0206024.
- [KS90] P.C. Kanellakis and S.A. Smolka. CCS expressions, finite state processes, and three problems of equivalence. *Information and Computation*, 86(1):43–68, 1990. doi:10.1016/0890-5401(90)90025-D.
- [Kul13] K. Kulakowski. Concurrent bisimulation algorithm. *ArXiv, CoRR*, abs/1311.7635, 2013.
- [LR94] I. Lee and S. Rajasekaran. A parallel algorithm for relational coarsest partition problems and its implementation. In D.L. Dill, editor, *Computer Aided Verification*, volume 818 of *Lecture Notes in Computer Science*, pages 404–414. Springer, 1994. doi:10.1007/3-540-58179-0_71.
- [MGH⁺21] J.J.M. Martens, J.F. Groote, L.B. van den Haak, H.P. Hijma, and A.J. Wijs. A linear parallel algorithm to compute bisimulation and relational coarsest partitions. In Gwen Salaün and Anton Wijs, editors, *Proc. FACS*, volume 13077 of *Lecture Notes in Computer Science*, pages 115–133. Springer, 2021. doi:10.1007/978-3-030-90636-8_7.
- [Mil80] R. Milner. *A Calculus of Communicating Systems*, volume 92 of *Lecture Notes in Computer Science*. Springer, 1980. doi:10.1007/3-540-10235-3.
- [Par81] D.M.R. Park. Concurrency and automata on infinite sequences. In P. Deussen, editor, *Proc. 5th GI-Conference on Theoretical Computer Science*, volume 104 of *Lecture Notes in Computer Science*, pages 167–183. Springer, 1981. doi:10.1007/BFb0017309.
- [PT87] R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–989, 1987. doi:10.1137/0216062.
- [PTB85] R. Paige, R.E. Tarjan, and R. Bonic. A linear time solution to the single function coarsest partition problem. *Theoretical Computer Science*, 40:67–84, 1985. doi:10.1016/0304-3975(85)90159-8.

- [RL98] S. Rajasekaran and I. Lee. Parallel algorithms for relational coarsest partition problems. *IEEE Transactions on Parallel and Distributed Systems*, 9(7):687–699, 1998. doi:10.1109/71.707548.
- [SV84] L.J. Stockmeyer and U. Vishkin. Simulation of parallel random access machines by circuits. *SIAM Journal of Computing*, 13(2):409–422, 1984. doi:10.1137/0213027.
- [WDMS20] T. Wißmann, U. Dorsch, S. Milius, and L. Schröder. Efficient and modular coalgebraic partition refinement. *Logical Methods Computer Science*, 16(1), 2020. doi:10.23638/LMCS-16(1:8)2020.
- [Wij15] A.J. Wijs. GPU accelerated strong and branching bisimilarity checking. In C. Baier and C. Tinelli, editors, *Proc. TACAS*, volume 9035 of *Lecture Notes in Computer Science*, pages 368–383. Springer, 2015. doi:10.1007/978-3-662-46681-0_29.