

Bisimulation by Partitioning is $\Omega((m+n) \log n)$

Jan Friso Groote   

Eindhoven University of Technology, The Netherlands

Jan Martens  

Eindhoven University of Technology, The Netherlands

Erik de Vink  

Eindhoven University of Technology, The Netherlands

Abstract

An asymptotic lowerbound of $\Omega((m+n) \log n)$ is established for partition refinement algorithms that decide bisimilarity on labeled transition systems. The lowerbound is obtained by subsequently analysing two families of deterministic transition systems — one with a growing action set and another with a fixed action set.

For deterministic transition systems with a one-letter action set, bisimilarity can be decided 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 the approach of Paige, Tarjan, and Bonic is not of help to develop a generic algorithm for deciding bisimilarity on labeled transition systems that is faster than the established lowerbound of $\Omega((m+n) \log n)$.

2012 ACM Subject Classification Theory of computation; Theory of computation → Interactive computation; Theory of computation → Design and analysis of algorithms

Keywords and phrases Bisimilarity, partition refinement, labeled transition system, lowerbound

Digital Object Identifier 10.4230/LIPIcs...

Funding Jan Martens: AVVA project NWO 612.001.751/TOP1.17.002

1 Introduction

Strong bisimulation [16, 13] is the gold standard for equivalence on labeled 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 [14] has a worst-case upperbound $O(m \log n)$ for establishing the bisimulation equivalence classes. Here, m is the number of transitions and n is the number of states 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 n by restricting the amount of work for refining blocks and moving states. Refining blocks is carried out by only investigating the smaller splitting blocks, 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 [8] and make the difference with the earlier $O(mn)$ algorithm of Kanellakis and Smolka [11].

The Paige-Tarjan algorithm, with its format of successive refinements of an initial partition till a fixpoint is reached, has been leading for variations and generalizations for deciding specific forms of (strong) bisimilarities, see e.g. [4, 6, 7, 18, 10]. We are interested in the question whether the Paige-Tarjan algorithm is computationally optimal. A result is provided by Berkholz et al. in a paper [2] that studies stable colourings of (coloured) graphs. More specifically, they show that for an undirected graph with n nodes and m edges, canonical coarsest bi-stable colouring is in $\Omega((m+n) \log n)$. Translated to LTSs, the result of [2] builds on the assumption that LTSs are essentially non-deterministic, i.e., every state has multiple outgoing transitions for the same label. A first contribution of the present paper is a lowerbound of the class of partition refinement algorithms



© J.F. Groote, J.J.M. Martens, E.P. de Vink, Eindhoven University of Technology;
licensed under Creative Commons License CC-BY 4.0

Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

XX:2 Bisimulation by partitioning is $\Omega((m+n)\log n)$

45 for deciding bisimilarity of deterministic LTSs. We define what a partition refinement algorithm is
46 and articulate the complexity in terms of the number of states that are moved. Then, a particular
47 family of (deterministic) LTSs, called bisplitters, is shown to require $n \log n$ work. This strengthens
48 the result of [2], actually answering an open question in it.

49 We obtain our lowerbound results assuming that algorithms use partition refinement. However, one
50 may wonder if a different approach than partition refinement can lead to a faster decision procedure
51 for bisimulation. For the specific case of deterministic LTSs with a singleton action set and a state
52 labelling, Robert Paige, Robert Tarjan and Robert Bonic propose a linear algorithm [15], which we
53 will refer to as Roberts' algorithm. In [5] it is proven that partition refinement à la Hopcroft has a
54 lowerbound of $\Omega(n \log n)$ in this case. Concretely, this means that Roberts' algorithm achieves the
55 essentially better performance by using a completely different technique than partition refinement to
56 determine the bisimulation equivalence classes.

57 Crucial for Roberts' algorithm is the ability to identify, in linear time, the bisimilarity classes of
58 cycles. In this paper we show that if the alphabet consists of at least two actions a rapid decision on
59 'cycles' as in [15] will not be of help to improve on the Paige-Tarjan algorithm for general LTSs. We
60 argue that the specialty in the algorithm of [15], viz. to be able to quickly decide the bisimilarity of
61 the states on cycles, can be captured by means of a stronger notion, namely an oracle, that provides
62 the bisimulation classes of the states of a so-called 'end structure', the counterpart in the multiple
63 action setting of a cycle in the single action setting. The oracle can be consulted to refine the initial
64 partition with respect to the bisimilarity on the end structures of the LTS for free. We show that for
65 the class of partition refinement algorithms enhanced with such an oracle, thus encompassing the
66 algorithm of [15], the $n \log n$ lowerbound persists for non-degenerate action sets.

67 The family of $n \log n$ -hard LTSs we use to establish the lowerbound, involve an action set of
68 $\log n$ actions. Building on the two results already mentioned, and exploiting ideas borrowed from [15]
69 to extend the bisimulation classes for the states in the end structures, i.e. cycles, to the states of
70 the complete LTS, we provide another family of (deterministic) LTSs that have only two actions.
71 Led by these LTSs we argue that for the two-action case the complexity of deciding bisimulation is
72 $\Omega((m+n)\log n)$, whether we use an oracle or not.

73 The document is structured as follows. In Section 2 we give the necessary preliminaries on
74 the problem. A recap of the linear algorithm of [15] is provided in Section 3. Next, we introduce
75 the family of deterministic LTSs \mathcal{B}_k for which we show in Section 4 that deciding bisimilarity
76 is $\Omega(n \log n)$ for the class of partition refinement algorithms and for which we establish in Section 5
77 an $\Omega(n \log n)$ lowerbound for the class of partition refinement algorithms enhanced with an oracle
78 for end structures. In Section 6 we introduce the family of deterministic LTSs \mathcal{C}_k , each involving
79 two actions only, to take the number of transitions m into account and establish an $\Omega((m+n)\log n)$
80 lowerbound for partition refinement with and without oracle for end structures. We wrap up with
81 concluding remarks.

82 2 Preliminaries

83 Given a set of states S , a *partition* of S is a set of sets of states $\pi \subseteq 2^S$ such that for all $B, B' \in \pi$
84 it holds that $B \neq \emptyset$, $B \cap B' = \emptyset$, and $\bigcup_{B \in \pi} B = S$. The elements of a partition are referred to
85 as blocks. A partition π of S induces an equivalence relation $=_\pi \subseteq S \times S$, where for two states
86 $s, t \in S$, $s =_\pi t$ iff the states are in the same block, i.e. there is a block $B \in \pi$ such that $s, t \in B$.
87 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
88 $B' \in \pi'$ such that $B \subseteq B'$. It follows that each block of π' is the union of blocks of π . The refinement
89 is *strict* if $\pi \neq \pi'$. The common refinement of two partitions π and π' is the partition with blocks
90 $\{B \cap B' \mid B \in \pi, B' \in \pi'\}$. A sequence of partitions π_0, \dots, π_n is called a refinement sequence iff

91 π_{i+1} is a refinement of π_i , for all $0 \leq i < n$.

92 ► **Definition 1.** A labeled transition system with initial partition (LTS) $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ is given
 93 by a finite set of states S , a finite alphabet of actions \mathcal{A} , a transition relation $\rightarrow \subseteq S \times \mathcal{A} \times S$, and a
 94 partition π_0 of S . A labeled transition system with initial partition is called deterministic (dLTS) if
 95 the transition relation is a total function $S \times \mathcal{A} \rightarrow S$.

96 Note that we omit an initial state, as it is not relevant in this article. Note also that in the presence of
 97 an initial partition, an LTS with one action label represents a Kripke structure. For a dLTS with a set
 98 of states S and the initial partition $\pi_0 = \{S\}$ we have that π_0 itself already represents bisimilarity,
 99 contrary to LTSs in general.

100 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
 101 of $(s, a, t) \in \rightarrow$. For dLTSs we occasionally write $L(s, a)$ for t , i.e., t is the image of the pair (s, a)
 102 of the function \rightarrow . We say that s reaches t via a iff $s \xrightarrow{a} t$. A state s reaches a set $U \subseteq S$ iff there is a
 103 state in U that is reached by s . A set of states $V \subseteq S$ is called *stable* under a set of states $U \subseteq S$ iff
 104 for all actions a either all states in V reach U via a , or no state in V reaches U via a . A partition π is
 105 stable under a set of states U iff each block $B \in \pi$ is stable under U . A partition π is called *stable* iff
 106 it is stable under all its blocks.

107 Following [16, 13], for an LTS L a symmetric relation $R \subseteq S \times S$ is called a bisimulation
 108 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 $t \xrightarrow{a} t'$
 109 for some $t' \in S$ such that $(s', t') \in R$. In the setting of the present paper, as we incorporate the
 110 initial partition in the definition of an LTS, bisimilarity is slightly non-standard. For a bisimulation
 111 relation R , we additionally require that it respects the initial partition π_0 of L , i.e. $(s, t) \in R$
 112 implies $s =_{\pi_0} t$. Two states $s, t \in S$ are called (strongly) bisimilar for L iff a bisimulation relation R
 113 exists with $(s, t) \in R$, notation $s \stackrel{\leftrightarrow}{L} t$. Bisimilarity is an equivalence relation on the set of states
 114 of L . We write $[s]_L^{\leftrightarrow}$ for the bisimulation equivalence class of the state s in L .

115 Partition refinement algorithms for deciding bisimilarity on LTSs start with an initial partition π_0 ,
 116 which is then repeatedly refined until a stable partition is reached. This stable partition is then the
 117 coarsest stable partition of the LTS refining π_0 and coincides with bisimilarity [11, 14].


118 We define that an algorithm is a partition refinement algorithm if it constructs a valid sequence of
 119 partitions. Concretely this means that a state s in a block B is only assigned to another block if there
 120 is reason to do so; there is a splitter block B' in the current partition to which s has a transition and
 121 some other state in B does not, or the other way around. Thus, the block B is not stable under the
 122 block B' . **Moreover, states that cannot be distinguished based on the current partition, i.e. there is
 123 no such block B' such that only one of the states reaches B' , cannot be a part of different blocks. (I
 124 think obsolete)** Finally, we insist that each subsequent partition reflects some progress, i.e., π_{i+1} is a
 125 strict refinement of π_i . This leads to the following notion of a valid refinement and a valid partition
 126 sequence.

127 ► **Definition 2.** Let $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ be an LTS, and π a partition of S . We call a refinement π'
 128 of π a valid refinement with respect to L , if the following criteria hold.

- 129 (a) π' is a strict refinement of π ;
 130 (b) if $s \neq_{\pi'} t$ for $s, t \in S$, then (i) $s \neq_{\pi} t$ or (ii) $s' \in S$ exist such that for some $a \in \mathcal{A}$, $s \xrightarrow{a} s'$ and
 131 for all $t' \in S$ such that $t \xrightarrow{a} t'$, we have $s' \neq_{\pi} t'$ or vice versa.
 132 A sequence of partitions $\Pi = (\pi_0, \dots, \pi_n)$ is called *valid* iff every successive partition π_i , for
 133 $0 < i \leq n$, is a valid refinement of π_{i-1} , and, moreover, the partition π_n is stable.

134 When a partition π is refined into a partition π' , states that are in the same block but can reach different
 135 blocks can lead to a split of the block into smaller ones, each holding one of the states. This means

XX:4 Bisimulation by partitioning is $\Omega((m+n) \log n)$



■ **Figure 1** dLTS with one action label (not shown) and initial partition distinguishing states 1, 4, 7, 10, 12 from states 2, 3, 5, 6, 8, 9, 11, 13.

136 that a block $B \in \pi$ is split into k blocks $B_1, \dots, B_k \in \pi'$. The least amount of work is done for this
 137 operation, by creating new blocks for the least number of states possible. Thus, $B \in \pi$ is transformed
 138 into $B_1 \in \pi'$, say, the biggest block among B_1, \dots, B_k . Therefore, the so-called refinement cost rc
 139 of the refinement π' of π is given by

$$140 \quad rc(\pi, \pi') = \sum_{B \in \pi} |B| - \max_{B' \in \pi': B' \subseteq B} |B'|.$$

141 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
 142 an LTS L , we write $rc(L)$ for $\min\{rc(\Pi) \mid \Pi \text{ a valid refinement sequence for } L\}$. Note that this
 143 complexity measure is more general than the one used in [2]. In that complexity measure the total
 144 number of transitions between the splitting block and the witnesses block thereof is counted, which is
 145 always greater or equal than our refinement costs. this sentence


146 We characterise the states of LTSs by sequences of bits. The set of bits is denoted as $\mathbb{B} = \{0, 1\}$.
 147 Bit sequences of length up to and including k are written as $\mathbb{B}^{\leq k}$. The inverse of a bit b is denoted by \bar{b} .
 148 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'
 149 and write $\sigma \prec \sigma'$ if σ is a strict prefix of σ' . For a bit sequence $\sigma \in \mathbb{B}^k$, for any $i, j \leq k$ we write $\sigma[i]$
 150 to indicate the bit at position i starting from position 1. We write $\sigma[i:j] = \sigma[i]\sigma[i+1] \dots \sigma[j]$ to
 151 indicate the subword from position i to position j .

152 **3 Roberts' algorithm**

153 Most algorithms to determine bisimulation on an LTS use partition refinement. However, there is
 154 one notable exception. On the class of dLTSs with a singleton action alphabet, deciding the coarsest
 155 stable partition, i.e. bisimilarity, requires linear time only. This is due to the algorithm of Robert
 156 Paige, Robert Tarjan, and Robert Bonic [15], which we therefore aptly call Roberts' algorithm.

157 The algorithm exploits the structure of dLTSs with one label, of which examples are depicted in
 158 Figure 1 where the initial partition is indicated by single/double circles around the states. For each
 159 state in such a dLTS we can assign a unique cycle (also referred to as 'end structure' in the sequel),
 160 and that state is either on this cycle or following outgoing edges lead to that (unique) cycle. Below,
 161 we roughly sketch how Roberts' algorithm works. See [15] for details.

- 162 1. Build lassos and mark states to identify the cycles of the dLTS.
- 163 2. For each state there is exactly one word of specific bounded length of labels of the initial partitions
 164 of the states on the (unique) run from this state. this sentence Identify for each cycle the states
 165 with the lexicographic least such word. This can be done in linear time in the size of the cycle. If
 166 there are bisimilar states on the cycle, then the algorithm will identify them. States on different
 167 cycles can only be bisimilar if the number of non-bisimilar states on these cycles is the same. By



■ **Figure 2** Bisplitter \mathcal{B}_3 with initial partition $\{\{000, 001, 010, 011\}, \{100, 101, 110, 111\}\}$.

168 comparing cycles with the same number of bisimulation equivalence classes, starting with the
 169 lexicographic least state, it is then determined linearly for all states on final cycles whether they
 170 are bisimilar.

171 **3.** By a backward calculation along the paths leading to the cycles the bisimilarity equivalence
 172 classes for the states not on cycles can then be determined in linear time as well.

173 A striking observation is that any partition refinement algorithm, using a valid sequence of partitions,
 174 requires a refinement cost of $\Omega(n \log n)$ to calculate which states are bisimilar for the dLTSs to which
 175 Roberts' algorithm applies. This follows from results in [3, 5] where it is shown that Hopcroft's
 176 algorithm [8] cannot have a better running time than $\Omega(n \log n)$. Below we come back to this
 177 observation, showing that the ideas in Roberts' algorithm cannot be exploited to come up with a linear
 178 algorithm for bisimulation if the LTS is either nondeterministic, or has more than one action label.

179 **4** \mathcal{B}_k is $\Omega(n \log n)$ for partition refinement

180 In this section we introduce a family of deterministic LTSs called bisplitters on which the cost of any
 181 partition refinement algorithm is $\Omega(n \log n)$ where n is the number of states. With some modification
 182 we obtain in Section 6 a family of LTSs that has the bound $\Omega((n + m) \log n)$ where m is the number
 183 of transitions.

184 ► **Definition 3.** For $k > 1$, the bisplitter $\mathcal{B}_k = (S, \mathcal{A}_k, \rightarrow, \pi_0)$ is defined as the dLTS that has the
 185 set $S = \{\sigma \mid \sigma \in \mathbb{B}^k\}$ as its set of states, the set $\mathcal{A}_k = \{a_1, \dots, a_{k-1}\}$ as its set of actions, the
 186 relation

$$187 \quad \{\sigma \xrightarrow{a_i} \sigma \mid \sigma \in S, 1 \leq i < k: \sigma[i+1] = 0\} \cup$$

$$188 \quad \{\sigma \xrightarrow{a_i} \sigma[1:i-1]\overline{\sigma[i]}0^{k-i} \mid \sigma \in S, 1 \leq i < k: \sigma[i+1] = 1\}$$

190 as its transition function, and the set $\pi_0 = \{\{\sigma \in S \mid \sigma[1] = 0\}, \{\sigma \in S \mid \sigma[1] = 1\}\}$ as its
 191 initial partition.

192 Thus the bisplitter \mathcal{B}_k has 2^k states, viz. the bitstrings of length k , and \mathcal{B}_k has $k-1$ action labels. It
 193 has $(k-1)2^k$ transitions: (i) a self-loop for bitstring σ with label a_i if the $i+1$ -th bit of σ equals 0;

XX:6 Bisimulation by partitioning is $\Omega((m+n)\log n)$

194 (ii) otherwise, i.e. when $i+1$ -th bit of σ equals 1, the bitstring σ has for label a_i a transition to the
 195 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.
 196 The initial partition π_0 distinguishes the bitstrings starting with 0 from those starting with 1. A
 197 drawing of bisplitter \mathcal{B}_3 is given in Figure 2. We see, e.g., for the bitstring $\sigma = 101$ an a_1 -transition
 198 to itself, as $\sigma[2] = 0$, and an a_2 -transition to 110 , as $\sigma[3] = 1$.

199 ► **Definition 4.** For any string $\sigma \in \mathbb{B}^{\leq k}$, we define the prefix block B_σ of \mathcal{B}_k to be the block
 200 $B_\sigma = \{\sigma' \in \mathbb{B}^k \mid \sigma \preceq \sigma'\}$.

201 The following lemma collects a number of results related to prefix blocks.

202 ► **Lemma 5.** Let $k \geq 2$ and let the dLTS $\mathcal{B}_k = (\mathbb{B}^k, \mathcal{A}_k, \rightarrow, \pi_0)$ be the k -th bisplitter. Let the
 203 sequence $\Pi = (\pi_0, \dots, \pi_n)$ be a valid refinement sequence. Then it holds that

- 204 (a) Every partition π_i of Π contains prefix blocks only.
 205 (b) If partition π_i of Π contains a prefix block B_σ with $|\sigma| < k$, then π_i is not stable.
 206 (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 $B_{\sigma 0} \in \pi_{i+1}$.

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

209 Assume, to arrive at a contradiction, that there is a block $B \in \pi_{i+1}$ that is not a prefix block.
 210 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
 211 that σ is a common prefix of all elements of B . We can choose θ such that $\sigma\theta$ is the longest common
 212 prefix of all elements of B . Since every singleton of \mathbb{B}^k is a prefix block, B is not a singleton. This
 213 means that $|\sigma\theta| < k$, and that there are some elements σ_1 and σ_2 of B such that $\sigma\theta 0$ is a prefix of σ_1
 214 and $\sigma\theta 1$ is a prefix of σ_2 . Because B is not a prefix block, there must exist at least one $\tau \in \mathbb{B}^k$ with a
 215 prefix $\sigma\theta$ such that $\tau \notin B$. Obviously, we have either (i) $\sigma\theta 0$ is a prefix of τ , or (ii) $\sigma\theta 1$ is a prefix
 216 of τ . We will show that in both these cases τ in fact belongs to B in π_{i+1} , which is a contradiction.
 217 This also means that for any $B \in \pi_{i+1}$, where $B \subseteq B_\sigma$ for some prefix block $B_\sigma \in \pi_i$, we have that
 218 B is a prefix block for the prefix $\sigma\theta$ (i.e. $B = B_{\sigma\theta}$) where $\sigma\theta$ is the longest common prefix of all
 219 elements of B .

220 (i) Suppose $\sigma\theta 0$ is a prefix of τ . We will show that τ and σ_1 belong to the same block in π_{i+1}
 221 because for each a_j (where $1 \leq j < k$) the states σ'_1 and τ' , such that $\sigma_1 \xrightarrow{a_j} \sigma'_1$ and $\tau \xrightarrow{a_j} \tau'$,
 222 belong to the same block in π_i . There are three cases:

- 223 – $j < |\sigma\theta|$: Since $\sigma\theta$ is a prefix of both σ_1 and τ , we have $\sigma_1[j+1] = \tau[j+1]$.
 224 – If $\sigma_1[j+1] = \tau[j+1] = 0$, then $\sigma'_1 = \sigma_1$ and $\tau' = \tau$. Obviously, both σ'_1 and τ' belong
 225 to B_σ (since σ_1 and τ belong to B_σ).
 226 – If $\sigma_1[j+1] = \tau[j+1] = 1$, then both σ'_1 and τ' are of the form $\rho[1 : j-1]\overline{\rho[j]}0^{k-j}$
 227 where $\rho = \sigma\theta$, and we have $\sigma'_1 = \tau'$, so they clearly belong to the same block of π_i .
 228 – $j = |\sigma\theta|$: Since $\sigma_1[j+1] = \tau[j+1] = 0$, we have $\sigma'_1 = \sigma_1$ and $\tau' = \tau$, and obviously both
 229 σ'_1 and τ' belong to B_σ (since σ_1 and τ belong to B_σ).
 230 – $j > |\sigma\theta|$: In fact, for arbitrary ρ , performing a_j with $j > |\sigma\theta|$ in $\sigma\theta\rho$ leads to $\sigma\theta\rho'$ (for both
 231 cases, where $(\sigma\theta\rho)[j+1]$ is 0 or 1). In particular this means that if $j > |\sigma\theta|$ and $\sigma_1 \xrightarrow{a_j} \sigma'_1$
 232 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 .

233 (ii) Now, suppose $\sigma\theta 1$ is a prefix of τ . We will show that τ and σ_2 belong to the same block in π_{i+1}
 234 because for each a_j (where $1 \leq j < k$) the states σ'_2 and τ' , such that $\sigma_2 \xrightarrow{a_j} \sigma'_2$ and $\tau \xrightarrow{a_j} \tau'$,
 235 belong to the same block in π_i . There are three cases:

- 236 – $j < |\sigma\theta|$: Similar as in (i).
 237 – $j = |\sigma\theta|$: Since $\sigma_1[j+1] = \tau[j+1] = 1$, we have $\sigma'_1 = \tau' = \rho[1 : j-1]\overline{\rho[j]}0^{k-1}$ where
 238 $\rho = \sigma\theta$, so clearly σ'_1 and τ' are in a same block in π_i .

239 – $j > |\sigma\theta|$: Similar as in (i).

240

241 **Proof (b).** Suppose $B_\sigma \in \pi_i$ and $|\sigma| = \ell < k$. Let $\theta \in \mathbb{B}^*$ be such that $\sigma_1 = \sigma\theta\theta$ and $\sigma_2 = \sigma1\theta$.
 242 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
 243 hence π_i isn't either. ◀

244 **Proof (c).** We show that for a prefix block $B_\sigma \in \pi_i$, a bit $b \in \mathbb{B}$ and all $\theta, \theta' \in \mathbb{B}^{k-(|\sigma|+1)}$ the states
 245 $\sigma_1 = \sigma b\theta$ and $\sigma_2 = \sigma b\theta'$ are not split by any action a_j , for $1 \leq j < k$, and thus are in the same
 246 block of π_{i+1} . Pick j , $1 \leq j < k$, and suppose $\sigma_1 \xrightarrow{a_j} \sigma'_1$, $\sigma_2 \xrightarrow{a_j} \sigma'_2$, i.e., $\sigma'_1 = \mathcal{B}_k(\sigma_1, a_j)$ and
 247 $\sigma'_2 = \mathcal{B}_k(\sigma_2, a_j)$. If $j \leq |\sigma|$ and $\sigma[j] = 0$ then $\sigma'_1 = \sigma_1$ and $\sigma'_2 = \sigma_2$ hence both $\sigma'_1, \sigma'_2 \in B_\sigma$ don't
 248 split for a_j . If $j \leq |\sigma|$ and $\sigma[j] = 1$ then $\sigma'_1 = \sigma'_2$ and don't split for a_j either. If $j > |\sigma|$ then both
 249 $\sigma'_1, \sigma'_2 \in B_\sigma$ and don't split for a_j either. ◀

250 With the help of the above lemma, clarifying the form of the partitions in a valid refinement sequence
 251 for the bisplitter family, we are able to obtain a lowerbound for any partition refinement algorithm
 252 acting on it.

253 ▶ **Theorem 6.** For any $k > 1$, application of partition refinement to the bisplitter \mathcal{B}_k has refinement
 254 costs $rc(\mathcal{B}_k) \in \Omega(n \log n)$ where $n = 2^k$ is the number of states of \mathcal{B}_k .

255 **Proof.** Let $\Pi = \pi_0, \dots, \pi_n$ be a valid refinement sequence for \mathcal{B}_k . By items a and b of Lemma 5,
 256 we have $\pi_n = \{ \{s\} \mid s \in \mathbb{B}^k \}$ since π_n is stable. Item c of Lemma 5 implies that in every refinement
 257 step (π_i, π_{i+1}) a block is kept or it is refined in two blocks of equal size. The cost of refining the
 258 block B_σ , for $|\sigma| < k$, into B_{σ_0} and B_{σ_1} is the number of states in B_{σ_0} or B_{σ_1} , which are the
 259 same and equal to $\frac{1}{2}2^{k-|\sigma|}$. Therefore, we have $rc(\mathcal{B}_k, \Pi) = \sum_{\ell=1}^{k-1} 2^\ell \frac{1}{2} 2^{k-\ell} = \sum_{\ell=1}^{k-1} \frac{1}{2} 2^k =$
 260 $(k-1)2^{k-1}$. If n is the number of states of \mathcal{B}_k , it holds that $n = 2^k$, thus $k-1 = \log \frac{1}{2}n$. Hence,
 261 $rc(\mathcal{B}_k, \Pi) = \frac{1}{2}n \log \frac{1}{2}n$ which is in $\Omega(n \log n)$. ◀

262 **5** \mathcal{B}_k is $\Omega(n \log n)$ for partition refinement with an oracle

263 One may wonder whether the approach of calculating bisimulation equivalence classes will work on
 264 transition systems with non-degenerate action sets as well as the Roberts' algorithm guarantees a
 265 linear performance for the degenerate case. In order to capture the approach of [15], we augment the
 266 class of partition refinement algorithms with an oracle. At the start of the algorithm the oracle can be
 267 consulted to identify the bisimulation classes for designated states, viz. for those that are in an 'end
 268 structure', the counterpart of the cycles in [15]. This results in a refinement of the initial partition;
 269 partition refinement then starts from the updated partition.

270 Thus, we can ask the oracle to provide the bisimulation classes of all elements in an end structure
 271 of the LTS at hand. This yields a new partition, viz. the common refinement of the initial partition,
 272 on the one hand, and the partition induced by the bisimulation equivalence classes as given by the
 273 oracle and the complement of their union, on the other hand. The work that remains to be done is
 274 establishing the bisimulation equivalence classes, with respect to the initial partition, for the states
 275 not in any end structure. We will establish that a partition refinement algorithm strengthened with
 276 such an oracle will not improve upon partition refinement.

277 We first define the notion of an end structure of an LTS and the associated notion of an end structure
 278 partition.

XX:8 Bisimulation by partitioning is $\Omega((m+n)\log n)$

279 ► **Definition 7.** Given an LTS $L = (S, \mathcal{A}, \rightarrow, \pi_0)$, a non-empty subset $S' \subseteq S$ is called an end
 280 structure of L , if S' is a minimal set of states closed under all transitions. Moreover, $es(L) = \{ S' \subseteq$
 281 $S \mid S' \text{ end structure of } L \}$ and $\pi_{es} = \{ [s]_L^{\leftrightarrow} \mid s \in \bigcup es(L) \} \cup \{ S \setminus \{ [s]_L^{\leftrightarrow} \mid s \in \bigcup es(L) \} \} \setminus \{ \emptyset \}$ is
 282 called the end structure partition of L .

283 Like the cycles of [15], an LTS can have multiple end structures. The end structure partition π_{es}
 284 consists of the bisimilarity equivalence classes of L containing a state of an end structure, completed
 285 with a block of the remaining states that are not in an end structure, and not bisimilar to any state in
 286 an end structure (if not empty).

287 ► **Lemma 8.** Let $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ be a dLTS.

288 (a) If $|\mathcal{A}| = 1$ then $es(L)$ consists of all cycles in L .

289 (b) Every $s \in S$ has a path to an end structure of L .

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

292 (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 U is
 293 closed under all transitions. The minimal subset $U' \subseteq U$ which is still closed under all transitions is
 294 an end structure of L and reachable by s . ◀

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

299 ► **Definition 9.** A partition refinement algorithm with end structure oracle yields for an LTS
 300 $L = (S, \mathcal{A}, \rightarrow, \pi_0)$ a valid refinement sequence $\Pi = (\pi'_0, \pi_1, \dots, \pi_n)$ where π'_0 is the common
 301 refinement of the initial partition π_0 and the end structure partition π_{es} of L . The partition π'_0 is
 302 called the updated initial partition of L .

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

309 ► **Lemma 10.** For the bisplitter $\mathcal{B}_k = (S, \mathcal{A}, \rightarrow, \pi_0)$, for some $k > 2$, let π'_0 be the updated initial
 310 partition. Then every valid refinement sequence $\Pi = (\pi'_0, \pi_2, \dots, \pi_n)$ for the updated bisplitter
 311 $\mathcal{B}'_k = (S, \mathcal{A}, \rightarrow, \pi'_0)$ satisfies $rc(\Pi) \geq rc(\mathcal{B}_{k-2})$.

312 **Proof.** Observe that there are only two end structures in \mathcal{B}_k , viz. the singletons of the two states
 313 with 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
 314 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 $b \in \mathbb{B}$, $j \geq 0$
 315 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
 316 $\sigma \xrightarrow{a_{j+1}} b0^{j-1}10^{k-(j+1)}$ while $b0^{j-1}10^{k-(j+1)}$ reaches 0^k or 10^{k-1} by induction.

317 By Lemma 5, every state $\sigma \in \mathbb{B}^k$ of \mathcal{B}_k is in its own bisimulation equivalence class $\{\sigma\}$. It
 318 follows that the updated initial partition π'_0 equals $\{ \{0^k\}, \{10^{k-1}\}, B_0 \setminus \{0^k\}, B_1 \setminus \{10^{k-1}\} \}$. We
 319 claim that if a sequence $\Pi = (\pi'_0, \pi_1, \dots, \pi_n)$ for \mathcal{B}_k exists with costs $rc(\Pi) \leq rc(\mathcal{B}_{k-2})$, then
 320 also a valid refinement sequence Π' for \mathcal{B}_{k-2} exists with costs smaller than $rc(\mathcal{B}_{k-2})$ which yields
 321 a contradiction, since, by definition, $rc(\mathcal{B}_{k-2})$ are the minimum costs over all valid refinement
 322 sequence for \mathcal{B}_{k-2} .

323 So, assume $\Pi = (\pi'_0, \pi_1, \dots, \pi_n)$ is a valid refinement sequence for \mathcal{B}_k and $rc(\Pi) \leq rc(\mathcal{B}_{k-2})$.
 324 We obtain a valid refinement sequence Π' for \mathcal{B}_{k-2} in two steps. First, we use the projection
 325 function p from partitions on \mathbb{B}^k to partitions on \mathbb{B}_{k-2} that removes the prefix 11 from strings in a
 326 block (or ignores the block if such string is absent), i.e. $p(\pi) = \{ \{ \sigma \mid 11\sigma \in B \} \mid B \in \pi \} \setminus \{ \emptyset \}$. In
 327 particular, $p(\pi_0) = \{ \{ \sigma \mid \sigma \in \mathbb{B}^{k-2} \} \}$. Second, the function P from refinement sequences of \mathcal{B}_k to
 328 refinement sequences of \mathcal{B}_{k-2} , removes, in addition to application of p to each constituent partition,
 329 duplicate partitions from the sequence. Then $\Pi' = P(\Pi)$, say $\Pi' = (\varrho_0, \varrho_1, \dots, \varrho_\ell)$.

330 We have $\varrho_0 = p(\pi'_0) = \{B_\varepsilon\}$. Next we observe that $\varrho_1 = \pi_0^{k-2} = \{B_0, B_1\}$ the initial partition
 331 of \mathcal{B}_{k-2} , containing the prefix blocks of 0 and 1: Take any two different states $b\theta, b\theta' \in \mathbb{B}^{k-2}$,
 332 for a bit $b \in \mathbb{B}$ and strings $\theta, \theta' \in \mathbb{B}^{k-3}$ that are not in the same block of ϱ_1 . Let $i, 0 \leq i < n$
 333 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
 334 refining π_i into π_{i+1} . But no action a_j witnesses such a split: (i) $\mathcal{B}_k(11b\theta, a_1) = \mathcal{B}_k(11b\theta', a_1)$
 335 as both equal 0^k ; (ii) $\mathcal{B}_k(110\theta, a_2) = 110\theta \in B_1 \setminus \{10^{k-1}\}$ and $\mathcal{B}_k(110\theta', a_2) = 110\theta' \in$
 336 $B_1 \setminus \{10^{k-1}\}$; (iii) $\mathcal{B}_k(111\theta, a_2) = \mathcal{B}_k(111\theta', a_2)$, viz. are equal to 10^{k-1} ; (iv) for $j > 2$ it holds
 337 that $\mathcal{B}_k(11b\theta, a_j), \mathcal{B}_k(11b\theta', a_j) \in B_1 \setminus \{10^{k-1}\}$. Since $\varrho_1 \neq \varrho_0$, ϱ_1 has at least two blocks.
 338 Hence these must be B_0 and B_1 .

339 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
 340 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 $\mathcal{B}_k(11\sigma, a_{j+2}) =$
 341 $11\sigma'$, a direct consequence of the definition of the transition functions of \mathcal{B}_{k-2} and \mathcal{B}_k . From this
 342 we obtain

$$343 \quad \sigma =_{\varrho_i} \sigma' \iff 11\sigma =_{\pi_h} 11\sigma' \tag{1}$$

344 provided $\varrho_i = p(\pi_h)$, for $0 \leq i \leq \ell$, via the definition of the projection function p . Now, consider
 345 subsequent partitions ϱ_i and ϱ_{i+1} in Π' . Let $h, 0 \leq h < n$, be such that $\varrho_i = p(\pi_h)$ and $\varrho_{i+1} = p(\pi_{h+1})$.
 346 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 $\alpha \in I$,
 347 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
 348 of the refinement of ϱ_i into ϱ_{i+1} is justified by the validity of π_{h+1} into π_h . If $\sigma =_{\varrho_i} \sigma'$ and
 349 $\sigma \neq_{\varrho_{i+1}} \sigma'$ for $\sigma, \sigma' \in \mathbb{B}^{k-2}$, then $\sigma, \sigma' \in B_0$ or $\sigma, \sigma' \in B_1$ since ϱ_i is a refinement of ϱ_0 . Moreover,
 350 $11\sigma =_{\pi_h} 11\sigma'$ and $11\sigma \neq_{\pi_{h+1}} 11\sigma'$ by (1). Hence, by validity, $\mathcal{B}_k(11\sigma, a_j) \neq_{\pi_h} \mathcal{B}_k(11\sigma', a_j)$.
 351 Clearly $j \neq 1$. Also, $j \neq 2$, since $\sigma[1] = \sigma'[1]$ we have $(11\sigma)[3] = (11\sigma')[3]$. Therefore,
 352 $\mathcal{B}_{k-2}(\sigma, a_{j-2}) \neq_{\pi_h} \mathcal{B}_{k-2}(\sigma', a_{j-2})$, showing the refinement of ϱ_i into ϱ_{i+1} to be valid.

353 Finally, since every block in π_n is a singleton, this is also the case for ϱ_ℓ . Thus, ϱ_ℓ is indeed the
 354 coarsest partition as required for Π' to be a valid refinement sequence for \mathcal{B}_{k-2} . Every refinement
 355 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.
 356 Therefore, since $\varrho_i = p(\pi_i)$ and $\varrho_{i+1} = p(\pi_{i+1})$, we have $rc(\varrho_i, \varrho_{i+1}) \leq rc(\pi_h, \pi_{h+1})$, and hence
 357 $rc(\Pi) = \sum_{h=1}^n rc(\pi_{h-1}, \pi_h) \geq \sum_{i=1}^\ell rc(\varrho_{i-1}, \varrho_i) = rc(\Pi') \geq rc(\mathcal{B}_{k-2})$, as was to be shown. \blacktriangleleft

358 Next we combine the above lemma with the lowerbound provided by Theorem 6 in order to prove the
 359 main result of this section.

360 **► Theorem 11.** *Any partition refinement algorithm with end structure oracle to decide bisimilarity*
 361 *for a dLTS is $\Omega(n \log n)$.*

362 **Proof.** Let \mathcal{B}'_k be the updated bisplitter (with the initial partition π'_0 containing $\{0^k\}, B_0 \setminus \{0^k\}$,
 363 $\{10^{k-1}\}$, and $B_1 \setminus \{10^{k-1}\}$ as given by the oracle for end structures rather than the partition π_0
 364 containing B_0 and B_1). By Lemma 10 we have, for $k > 2$, that $rc(\mathcal{B}'_k) \geq rc(\mathcal{B}_{k-2})$. By Theorem 6
 365 we know that $rc(\mathcal{B}_{k-2}) \geq \frac{1}{2}n' \log \frac{1}{2}n'$ where $n' = 2^{k-2}$ is the number of states of \mathcal{B}_{k-2} . It holds
 366 that $n' = \frac{2^{k-2}}{2^k}n = \frac{1}{4}n$. So $rc(\mathcal{B}'_k) \geq \frac{1}{8}n \log \frac{1}{8}n$ from which we conclude that deciding bisimilarity
 367 for \mathcal{B}_k with the help of an oracle for the end structures is $\Omega(n \log n)$. \blacktriangleleft

368 **6** \mathcal{C}_k is $\Omega((m+n)\log n)$ for partition refinement

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

374 To this end we adapt the construction of \mathcal{B}_k at two places. Given an action alphabet \mathcal{A} of \mathcal{B}_k of
 375 $k-1$ actions, we introduce for each $\sigma \in \mathbb{B}^k$, a stake of 2^k states. Moreover, for each stake we add a
 376 tree gadget. These gadgets have height $\lceil \log(\frac{k-1}{2}) \rceil$ to accommodate $\lceil (k-1)/2 \rceil$ leaves.

377 **► Definition 12.** Let $k > 1$, \mathcal{B}_k be the k -th bisplitter, and $\mathbb{A} = \{a, b\}$. The dLTS $\mathcal{C}_k =$
 378 $(S_k^{\mathcal{C}}, \mathbb{A}, \rightarrow_{\mathcal{C}}, \pi_0^{\mathcal{C}})$, over the action set \mathbb{A} ,

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

$$380 \quad S_k^{\mathcal{C}} = \{ [\sigma, \ell] \in \mathbb{B}^k \times \mathbb{N} \mid 1 \leq \ell \leq 2^k \} \cup \\ \{ \langle \sigma, w \rangle \in \mathbb{B}^k \times \mathbb{A}^* \mid 0 \leq |w| \leq \lceil \log(\frac{k-1}{2}) \rceil \},$$

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

$$382 \quad \begin{array}{lll} [\sigma, \ell] & \xrightarrow{\alpha}_{\mathcal{C}} & [\sigma, \ell + 1] \quad \text{for } \sigma \in \mathbb{B}^k, 1 \leq \ell \leq 2^k, \alpha \in \mathbb{A} \\ [\sigma, 2^k] & \xrightarrow{\alpha}_{\mathcal{C}} & \langle \sigma, \varepsilon \rangle \quad \text{for } \sigma \in \mathbb{B}^k, \alpha \in \mathbb{A} \\ \langle \sigma, w \rangle & \xrightarrow{\alpha}_{\mathcal{C}} & \langle \sigma, w\alpha \rangle \quad \text{for } \sigma \in \mathbb{B}^k, |w| < \lceil \log(\frac{k-1}{2}) \rceil, \alpha \in \mathbb{A} \\ \langle \sigma, w \rangle & \xrightarrow{\alpha}_{\mathcal{C}} & [\sigma', 1] \quad \text{for } \sigma \in \mathbb{B}^k, |w| = \lceil \log(\frac{k-1}{2}) \rceil, \text{bin}(w\alpha) = j, \mathcal{B}_k(\sigma, a_j) = \sigma', \end{array}$$

383 (c) and has the initial partition $\pi_0^{\mathcal{C}}$ defined as

$$384 \quad \pi_0^{\mathcal{C}} = \{ \{ [\sigma, \ell] \mid \sigma \in B_0 \}, \{ [\sigma, \ell] \mid \sigma \in B_1 \} \mid 1 \leq \ell \leq 2^k \} \cup \\ 385 \quad \{ \langle \sigma, w \rangle \in S_k^{\mathcal{C}} \mid \sigma \in \mathbb{B}^k, w \in \mathbb{A}^* \}.$$

386 The auxiliary function $\text{bin} : \mathbb{A}^{\leq \lceil \log(k-1) \rceil} \rightarrow \mathbb{N}$, used in item b is inductively defined by $\text{bin}(\varepsilon) = 0$,
 387 $\text{bin}(wa) = \min\{2 * \text{bin}(w), k-1\}$, and $\text{bin}(wb) = \min\{2 * \text{bin}(w)+1, k-1\}$.

388 We see that with each string $\sigma \in \mathbb{B}^k$ we associate in \mathcal{C}_k as many as 2^k stake states $[\sigma, 1], \dots, [\sigma, 2^k]$,
 389 one for each level ℓ , $1 \leq \ell \leq 2^k$. The stake states are traversed from the top $[\sigma, 1]$ to bot-
 390 tom $[\sigma, 2^k]$ on any string σ of length 2^k over \mathbb{A} . The tree gadget consists of a complete binary
 391 tree of height $\lceil \log(\frac{k-1}{2}) \rceil$ that hence has at least $\lceil (k-1)/2 \rceil$ leaves. Traversal down the tree takes a
 392 left child on action a , a right child on action b . Together with the two actions of \mathbb{A} , $k-1$ source-label
 393 pairs can be encoded. 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
 394 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
 395 string $w\alpha$ for $\alpha \in \mathbb{A}$ is the binary encoding according of bin of the index j . An α -transition thus
 396 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
 397 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
 398 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
 399 in C_1^{ℓ} , and the states of the tree gadgets collected in C_{ε} .

400 Figure 3 depicts the two-label layered 3-splitter \mathcal{C}_3 . Because also \mathcal{B}_k has an action set of size 2
 401 the tree gadgets only consist of the root node of the form $\langle \sigma, \varepsilon \rangle$. In Figure 2 of \mathcal{B}_3 we see that
 402 $101 \xrightarrow{a_1} 101$ and $101 \xrightarrow{a_2} 110$. In Figure 3 we have transitions $\langle 101, \varepsilon \rangle \xrightarrow{a} [101, 1]$ and
 403 $\langle 101, \varepsilon \rangle \xrightarrow{b} [110, 1]$ (dotted and dashed, respectively). Coloring of nodes is used to represent the
 404 initial partition $\pi_3^{\mathcal{C}}$ that separates 8 times, once for each level ℓ , the four states of the stakes in C_0^{ℓ} on
 405 the left from the four stake states in C_1^{ℓ} on the right, and the 8 tree states in C_{ε} at the bottom of the
 406 picture.

407 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
 408 corresponding outgoing transitions is drawn in Figure 4. The tree has height $\lceil \log((6-1)/2) \rceil =$

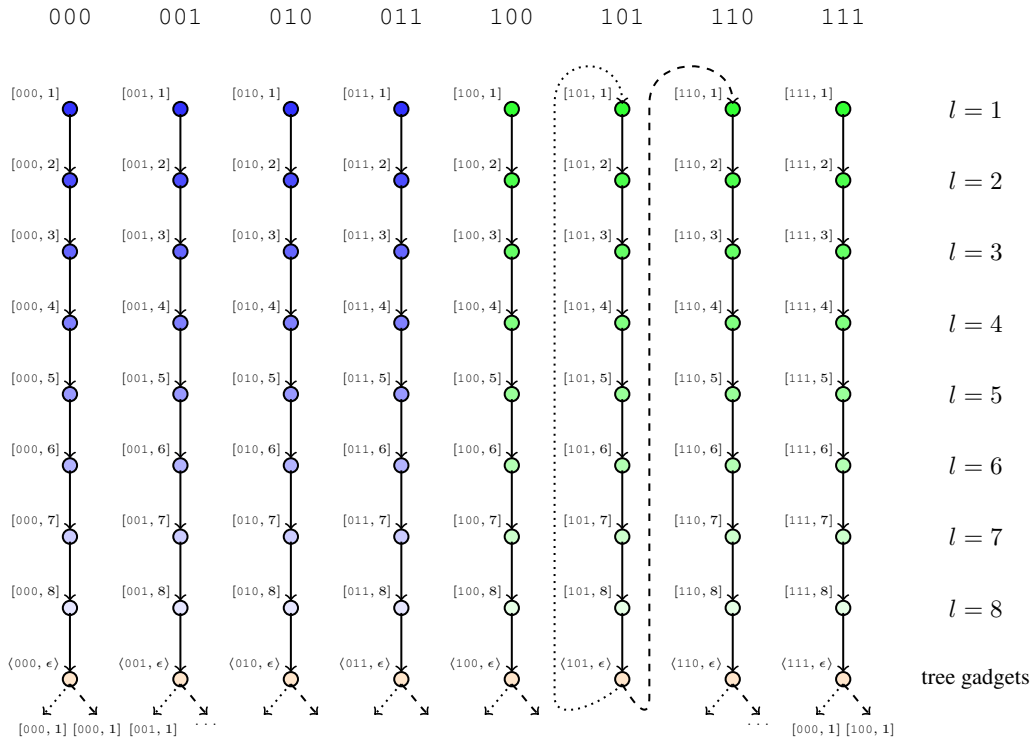


Figure 3 The partial layered bisplitter \mathcal{C}_3 with tree gadgets, the colors represent the initial partition




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

409 $\lceil \log \frac{5}{2} \rceil = 2$, hence it has $2^2 = 4$ leaves. Since each leaf has two outgoing transitions, one labeled a
 410 and one labeled b , the two leftmost leaves $\langle \sigma, aa \rangle$ and $\langle \sigma, ab \rangle$ are used with the two labels a and b to
 411 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
 412 transitions all simulating the a_5 -transition of σ .

413 The next lemma introduces three facts for the layered bisplitter \mathcal{C}_k that we need in the sequel. The
 414 first states that if two states in different stakes, but at the same level, are separated during partition
 415 refinement, then all corresponding states at lower levels are separated as well. The second fact helps
 416 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
 417 by a path from $[\sigma, 2^k]$ through the tree gadget of σ from root to leaf and then to the top state $[\sigma', 1]$
 418 of the stake of σ' . The word $w\alpha$ encountered going down and out the tree gadget corresponds to the
 419 action a_j according to the bin-function. Lastly, it is shown that no two pairs of different states within
 420 the stakes are bisimilar.

421 ► **Lemma 13.** *Let Π be a valid refinement sequence for \mathcal{C}_k and π a partition in Π .*

XX:12 Bisimulation by partitioning is $\Omega((m+n)\log n)$

422 (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
 423 $[\sigma, m], [\sigma', m] \in S$, for all levels m , $\ell \leq m \leq 2^k$, are in different blocks of π .

424 (b) If $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ are split in π , then exist $w \in \mathbb{A}^*$, $\alpha \in \mathbb{A}$, and $\sigma'_1, \sigma'_2 \in \mathbb{B}^k$ such that

$$425 \quad [\sigma_1, 2^k] \xrightarrow{w}_c^* \langle \sigma_1, w \rangle \xrightarrow{\alpha}_c [\sigma'_1, 1] \quad \text{and} \quad [\sigma_2, 2^k] \xrightarrow{w}_c^* \langle \sigma_2, w \rangle \xrightarrow{\alpha}_c [\sigma'_2, 1]$$

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

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

428 **Proof.** (a) For a proof by contradiction, suppose the partition π is the first partition of Π that falsifies
 429 the statement of the lemma. So $\pi \neq \pi_0^{\mathcal{C}}$, since for the initial refinement $\pi_0^{\mathcal{C}}$ the statement holds. Thus,
 430 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
 431 of π while the states $[\sigma, \ell+1], [\sigma', \ell+1]$ are in the same block of π and hence of π' . Since $[\sigma, \ell]$
 432 and $[\sigma', \ell]$ only have transitions to $[\sigma, \ell+1]$ and $[\sigma', \ell+1]$, respectively, that are in the same block π' ,
 433 the refinement wouldn't have been valid. We conclude that no falsifying partition π in Π exists and
 434 that the lemma holds.

435 (b) We first prove, by induction on $|w|$, that if $\langle \sigma_1, w \rangle$ and $\langle \sigma_2, w \rangle$ are split in π , then exist $w \in \mathbb{A}^*$
 436 and $\alpha \in \mathbb{A}$ such that $\langle \sigma_1, w \rangle \xrightarrow{v}_c^* \langle \sigma_1, wv \rangle \xrightarrow{\alpha}_c [\sigma'_1, 1]$ and $\langle \sigma_2, w \rangle \xrightarrow{v}_c^* \langle \sigma_2, wv \rangle \xrightarrow{\alpha}_c [\sigma'_2, 1]$ with
 437 $[\sigma'_1, 1]$ and $[\sigma'_2, 1]$ in different blocks of π . If w has maximal length, $|w| = \lceil \log(\frac{k-1}{2}) \rceil$ this is clear. If
 438 $\langle \sigma_1, w \rangle$ and $\langle \sigma_2, w \rangle$ are split, for $|w| < \lceil \log(k-1) \rceil - 1$, then either a -transitions or b -transitions lead
 439 to split states. By the induction hypothesis, suitable paths exists from the targets of such transitions.
 440 Adding the respective transition proves the induction hypothesis. Since $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ can only
 441 reach $\langle \sigma_1, \varepsilon \rangle$ and $\langle \sigma_2, \varepsilon \rangle$ the statement follows.

442 (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 $(\sigma_1, \sigma_2) \in R$
 443 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
 444 relation for \mathcal{B}_k . Note, that R respects $\pi_k^{\mathcal{B}}$, the initial partition of \mathcal{B}_k . Now, suppose $(\sigma_1, \sigma_2) \in R$
 445 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 $[\sigma_1, \ell] \xrightarrow{a_j}^*$
 446 $[\sigma_1, 2^k] \xrightarrow{a} \langle \sigma_1, \varepsilon \rangle \xrightarrow{w}_c^* \langle \sigma_1, w \rangle \xrightarrow{\alpha}_c [\sigma'_1, 1] \xrightarrow{a^{\ell-1}}^* [\sigma'_1, \ell]$ where $\text{bin}(w\alpha) = j$. Since $[\sigma_1, \ell]$ and
 447 $[\sigma_2, \ell]$ are bisimilar in \mathcal{C}_k , it follows that a corresponding path $[\sigma_2, \ell] \xrightarrow{a_j}^* [\sigma'_2, \ell]$ exists in \mathcal{C}_k with
 448 $[\sigma'_1, \ell]$ 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 $(\sigma'_1, \sigma'_2) \in R$.
 449 Hence, R is a bisimulation relation for \mathcal{B}_k indeed. Now, bisimilarity of \mathcal{B}_k is discrete. Thus, if two
 450 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 $\sigma_1 = \sigma_2$,
 451 and therefore $[\sigma_1, \ell] = [\sigma_2, \ell]$. ◀

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

454 ► **Lemma 14.** *It holds that $rc(\mathcal{C}_k) \geq 2^k rc(\mathcal{B}_k)$ for all $k > 1$.*

455 **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
 456 level ℓ , the sequence Π induces a valid refinement sequence Π^ℓ for \mathcal{B}_k .

457 The mapping p_ℓ assigns to a partition π of \mathcal{C}_k a partition $p_\ell(\pi)$ by putting

$$458 \quad p_\ell(\pi) = \{ \{ \sigma \in \mathbb{B}^k \mid [\sigma, \ell] \in B \} \mid B \in \pi \} \setminus \{ \emptyset \}.$$

459 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
 460 removing possible duplicates. We verify that Π^ℓ is a valid refinement sequence for \mathcal{B}_k .

461 First, we check that π_i^ℓ is a refinement of π_{i-1}^ℓ , for $1 \leq i \leq m$. Choose such an index i arbitrary.
 462 Let the index h with $1 \leq h \leq n$ by such that $p_\ell(\pi_{h-1}) = \pi_{i-1}^\ell$ and $p_\ell(\pi_i) = \pi_h^\ell$. For each
 463 block $B' \in \pi_i^\ell$ exists a block $B \in \pi_h$ such that $B' = p_\ell(B)$. Since π_h is a refinement of π_{h-1} ,

464 thus $B = \bigcup_r B_r$ for suitable $B_r \in \pi_h$. Note, $p_\ell(B_r) \in p_\ell(\pi_{h-1})$ for each index r . We have
 465 $B' = \bigcup_r p_\ell(B_r)$ with $p_\ell(B_r) \in \pi_{i-1}^\ell$, and π_i^ℓ is a refinement of π_{i-1}^ℓ .

466 Next, we verify that Π^ℓ is a valid refinement sequence for \mathcal{B}_k . Suppose the state $\sigma_1, \sigma \in S_k^{\mathcal{B}}$
 467 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
 468 refinement of a partition π_{h-1} into the partition π_h for a some index $h, 1 \leq h \leq n$. Then either
 469 (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) $\ell < 2^k$
 470 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 13
 471 that also $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$ are in different blocks of π_{h-1} . Thus, for the refinement of some π_{g-1}
 472 into $\pi_g, 1 \leq g \leq h \leq n$, splitted the two states $[\sigma_1, 2^k]$ and $[\sigma_2, 2^k]$. By Lemma 13 exist $w \in \mathbb{A}^*,$
 473 $\alpha \in \mathbb{A}$, and $\sigma'_1, \sigma'_2 \in \mathbb{B}^k$ such that

$$474 \quad [\sigma_1, 2^k] \xrightarrow{w}_c^* \langle \sigma_1, w \rangle \xrightarrow{\alpha}_c [\sigma'_1, 1] \quad \text{and} \quad [\sigma_2, 2^k] \xrightarrow{w}_c^* \langle \sigma_2, w \rangle \xrightarrow{\alpha}_c [\sigma'_2, 1]$$

475 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
 476 while $\sigma_1 \xrightarrow{a_j}_B \sigma'_1$ and $\sigma_2 \xrightarrow{a_j}_B \sigma'_2$ for $j = \text{bin}(w\alpha)$, which justifies splitting σ_1 and σ_2 for π_i^ℓ . We
 477 conclude that Π^ℓ is a valid refinement sequence for \mathcal{B}_k .

478 We have established that if Π is a valid refinement sequence for \mathcal{C}_k , then Π^ℓ is a valid refinement
 479 sequence for \mathcal{B}_k . The sequence Π^ℓ is obtained from Π by sifting out the blocks of Π 's partitions
 480 and removing repeated partitions. Therefore it holds that $rc(\Pi) \geq rc(\Pi^\ell)$. Since the mapping
 481 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
 482 $rc(\Pi) \geq \sum_{\ell=1}^{2^k} rc(\Pi^\ell) \geq 2^k rc(\mathcal{B}_k)$. Taking the minimum over all valid refinement sequences for \mathcal{C}_k
 483 we conclude that $rc(\mathcal{C}_k) \geq 2^k rc(\mathcal{B}_k)$ as was to be shown. \blacktriangleleft

484 With the above technical lemma in place, we are able to strengthen the $\Omega(n \log n)$ lowerbound of
 485 Theorem 6 to account for the number of transitions. The improved lowerbound is $\Omega((m+n) \log n)$,
 486 where m is the number of transitions and n the number of states.

487 **► Theorem 15.** *Deciding bisimilarity for dLTSs 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 dLTS.*

489 **Proof.** For the bisplitter \mathcal{B}_k , we know by Theorem 6 that $rc(\mathcal{B}_K) \geq 2^{k-1}(k-1)$. Thus, by
 490 Lemma 14, we obtain $rc(\mathcal{C}_K) \geq 2^{2k-1}(k-1)$. In the case of \mathcal{C}_k we have $n = 2^k(2^k + 2^{\lceil \log(k-1) \rceil} - 1)$
 491 and $m = 2n$. Hence $n+m \in \Theta(2^{2k-1})$ and $\log n \in \Theta(k-1)$, from which it follows that $rc(\mathcal{C}_k) \in$
 492 $\Omega((m+n) \log n)$. \blacktriangleleft

493 Underlying the proof of a lowerbound for deciding bisimilarity for the family of layered bisplitters \mathcal{C}_k
 494 is the observation that each \mathcal{C}_k can be seen as 2^k stacked instances of the ordinary bisplitters \mathcal{B}_k ,
 495 augmented with tree gadgets to handle transitions properly. The other essential ingredient for the
 496 proof of Theorem 15 is the complexity of deciding bisimilarity with a partition refinement algorithm
 497 on the \mathcal{B}_k family. The same reasoning applies when considering partition refinement algorithms with
 498 an oracle for end structures from Section 5. Also with an oracle the lowerbound of $\Omega((m+n) \log n)$
 499 remains.

500 **► Theorem 16.** *Any partition refinement algorithm with an oracle for end structures that decides*
 501 *bisimilarity for dLTSs is $\Omega((m+n) \log n)$.*

502 **Proof sketch.** The proof is similar to that of Lemma 10 and Theorem 15. Consider, for some $k > 2$,
 503 the layered bisplitter \mathcal{C}_k having initial partition π_0 . The dLTS \mathcal{C}_k has two end structures, viz. the
 504 set $S_0 \subset S_k^{\mathcal{C}}$ containing the states of the stake and accompanying tree gadget $S_0 = \{[0^k, \ell] \mid 1 \leq$
 505 $\ell \leq 2^k\} \cup \{\langle 0^k, w \rangle \mid w \in \mathbb{A}^*, |w| \leq \lceil \log(\frac{k-1}{2}) \rceil\}$ for 0^k and a similar $S_1 \subseteq S_k^{\mathcal{C}}$ for 10^{k-1} . The
 506 sets S_0 and S_1 are minimally closed under the transitions of \mathcal{C}_k . Other states, on the stake or tree

XX:14 Bisimulation by partitioning is $\Omega((m+n)\log n)$

507 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 .
508 The bisimulation classes S'_0 and S'_1 , say, with respect to $S_k^{\mathcal{C}}$ rather than π_0 , consist of S_0 and S_1
509 themselves plus a part of the tree gadgets for transitions in \mathcal{C}_k leading to S_0 and S_1 , respectively.

510 The update of the initial partition π_0 with oracle information, which concerns, ignoring the tree
511 gadgets, the common refinement of the layers on $\{[\sigma, \ell] \mid \sigma \in B_0\}$ and $\{[\sigma, \ell] \mid \sigma \in B_1\}$ on
512 the one hand, and (a trivial refinement of) the bisimulation classes S'_0 and S'_1 on the other hand, is
513 therefore equal to π on the stakes (and generally finer on the tree gadgets).

514 Then every valid refinement sequence $\Pi = (\pi'_0, \pi_2, \dots, \pi_n)$ for the updated dLTS $\mathcal{C}'_k =$
515 $(S, \mathcal{A}, \rightarrow, \pi'_0)$ satisfies $rc(\Pi) \geq rc(\mathcal{C}_{k-2})$. Following the lines of the proof of Lemma 10, we
516 can show that a valid refinement sequence Π for \mathcal{C}_k with updated initial partition π'_0 induces a valid
517 refinement sequence Π' for \mathcal{C}_{k-2} .

518 The number of states in \mathcal{C}_{k-2} is $\Theta(n)$ with n the number of states of \mathcal{C}_k , and number of transitions
519 in \mathcal{C}_{k-2} is $\Theta(m)$ with m the number of transitions of \mathcal{C}_k . Therefore, $rc(\Pi) \geq rc(\Pi') \geq rc(\mathcal{C}_{k-2})$,
520 from which we derive that any partition refinement algorithm with oracle for end structures involves
521 $\Theta((m+n)\log n)$ times moving a state for \mathcal{C}_k and that hence the algorithm is $\Omega((m+n)\log n)$. ◀

7 Conclusion

523 We have shown that, even when restricted to deterministic LTSs, it is not possible to construct an
524 algorithm based on partition refinement that is more efficient than $\Omega((m+n)\log n)$. This strengthens
525 the result of [2]. The bound obtained is preserved even when the algorithm is extended with an oracle
526 that can determine for specific states whether they are bisimilar or not in constant time. The oracle
527 proof technique enabled us to show that the algorithmic ideas underlying Roberts' algorithm [15] for
528 the one-letter alphabet case cannot be used to come up with a fundamentally faster enhanced partition
529 refinement algorithm for bisimulation. Of course, this is not addressing a generic lower bound to
530 decide bisimilarity on LTSs, nor proving the conjecture that the Paige-Tarjan algorithm is optimal
531 for deciding bisimilarity. It is conceivable that a more efficient algorithm exists that is not based on
532 partitioning for bisimulation. However, as it stands, no techniques are known to prove such a generic
533 algorithmic lowerbound, and all that do exist make assumptions on allowed operations, such as
534 the well-known lowerbound on sorting. Further investigations to obtain a more general lowerbound
535 may strengthen the oracle used even further, such that a wider range of algorithms is covered.

536 For the parallel setting, where deciding bisimilarity can be done faster indeed, a similar dichotomy
537 between the case of a single letter alphabet and of a multiple letter alphabet occurs. For LTSs with
538 multiple action labels a linear algorithm is proposed in [12], whereas for dLTSs with one action label
539 it is possible to calculate bisimulation in logarithmic time, cf. [9]. The question is raised in [17],
540 if a sub-linear parallel solution exists at all. Since this problem in a general setting is known to be
541 P-complete as shown in [1], it is generally believed that no logarithmic algorithm is possible. It
542 is worthwhile to transfer the results of this paper to a parallel setting, in order to better understand
543 whether it is possible to design parallel partition based algorithms for bisimulation on LTSs that have
544 a sub-linear complexity.

545 *Acknowledgements* We are grateful to the anonymous reviewers of CONCUR 2021 their thorough
546 reading and constructive feedback.

References

- 548 1 J. Balcázar, J. Gabarro, and M. Santha. Deciding bisimilarity is P-complete. *Formal Aspects of Computing*,
549 4(1):638–648, 1992.
- 550 2 C. Berkholz, P. Bonsma, and M. Grohe. Tight lower and upper bounds for the complexity of canonical
551 colour refinement. *Theory of Computing Systems*, 60(4):581–614, 2017.

- 552 **3** J. Berstel and O. Carton. On the complexity of Hopcroft’s state minimization algorithm. In M. Domaratzki
553 et al., editor, *Proc. CIAA 2004*, pages 35–44. LNCS 3317, 2004.
- 554 **4** P. Buchholz. Exact performance equivalence: An equivalence relation for stochastic automata. *Theoretical*
555 *Computer Science*, 215:263–287, 1999.
- 556 **5** G. Castiglione, A. Restivo, and M. Sciortino. Hopcroft’s algorithm and cyclic automata. In C. Martín-Vide
557 et al., editor, *Proc. LATA 2008*, pages 172–183. LNCS 5196, 2008.
- 558 **6** A. Dovier, C. Piazza, and A. Policriti. An efficient algorithm for computing bisimulation equivalence.
559 *Theoretical Computer Science*, 311:221–256, 2004.
- 560 **7** J.F. Groote, H.J. Rivera Verduzco, and E.P. de Vink. An efficient algorithm to determine probabilistic
561 bisimulation. *Algorithms*, 11(9):131,1–22, 2018.
- 562 **8** J. Hopcroft. An $n \log n$ algorithm for minimizing states in a finite automaton. In Z. Kohavi and A. Paz,
563 editors, *Theory of Machines and Computations*, pages 189–196. Academic Press, 1971.
- 564 **9** J. Jájá and Kwan Woo Ryu. An efficient parallel algorithm for the single function coarsest partition
565 problem. *Theoretical Computer Science*, 129(2):293–307, 1994.
- 566 **10** D.N. Jansen, J.F. Groote, J.J.A. Keiren, and A. Wijs. An $O(m \log n)$ algorithm for branching bisimilarity
567 on labelled transition systems. In A. Biere and D. Parker, editors, *Proc. TACAS*, pages 3–20. LNCS 12079,
568 2020.
- 569 **11** P.C. Kanellakis and S.A. Smolka. CCS expressions, finite state processes, and three problems of
570 equivalence. *Information and Computation*, 86(1):43–68, 1990.
- 571 **12** J. Martens, J.F. Groote, L. van de Haak, P. Hijma, and A. Wijs. A linear parallel algorithm to compute
572 bisimulation and relational coarsest partitions. In preparation.
- 573 **13** R. Milner. *A Calculus of Communicating Systems*. LNCS 92, 1980.
- 574 **14** R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM Journal on Computing*, 16(6):973–
575 989, 1987.
- 576 **15** R. Paige, R.E. Tarjan, and R. Bonic. A linear time solution to the single function coarsest partition
577 problem. *Theoretical Computer Science*, 40:67–84, 1985.
- 578 **16** D.M.R. Park. Concurrency and automata on infinite sequences. In P. Deussen, editor, *Proc. 5th GI-*
579 *Conference on Theoretical Computer Science*, pages 167–183. LNCS 104, 1981.
- 580 **17** S. Rajasekaran and I. Lee. Parallel algorithms for relational coarsest partition problems. *IEEE Transactions*
581 *on Parallel and Distributed Systems*, 9(7):687–699, 1998.
- 582 **18** T. Wißmann, U. Dorsch, S. Milius, and L. Schröder. Efficient and modular coalgebraic partition refinement.
583 *Logical Methods Computer Science*, 16(1), 2020.