

A More Efficient Simulation Algorithm on Kripke Structures

Francesco Ranzato

Dipartimento di Matematica, University of Padova, Italy

Abstract. A number of algorithms for computing the simulation preorder (and equivalence) on Kripke structures are available. Let Σ denote the state space, \rightarrow the transition relation and P_{sim} the partition of Σ induced by simulation equivalence. While some algorithms are designed to reach the best space bounds, whose dominating additive term is $|P_{\text{sim}}|^2$, other algorithms are devised to attain the best time complexity $O(|P_{\text{sim}}||\rightarrow|)$. We present a novel simulation algorithm which is both space and time efficient: it runs in $O(|P_{\text{sim}}|^2 \log |P_{\text{sim}}| + |\Sigma| \log |\Sigma|)$ space and $O(|P_{\text{sim}}||\rightarrow| \log |\Sigma|)$ time. Our simulation algorithm thus reaches the best space bounds while closely approaching the best time complexity.

1 Introduction

The simulation preorder is a fundamental behavioral relation widely used in process algebra for establishing system correctness and in model checking as a suitable abstraction for reducing the size of state spaces. The problem of efficiently computing the simulation preorder (and consequently simulation equivalence) on finite Kripke structures has been thoroughly investigated and generated a number of simulation algorithms. Both time and space complexities play an important role in simulation algorithms, since in several applications, especially in model checking, memory requirements may become a serious bottleneck as the input transition system grows.

Consider a finite Kripke structure where Σ denotes the state space, \rightarrow the transition relation and P_{sim} the partition of Σ induced by simulation equivalence. The best simulation algorithms are those by, in chronological order, Gentilini, Piazza and Policriti (GPP) [3] (subsequently corrected in [4]), Ranzato and Tapparo (RT) [10,12], Markovski (Mar) [8], Cécé (Space-Céc and Time-Céc) [2]. The simulation algorithms GPP and RT are designed for Kripke structures, while Space-Céc, Time-Céc and Mar are for more general labeled transition systems. Their space and time complexities are summarized in the following table.

Algorithm	Space complexity	Time complexity
Space-Céc [2]	$O(P_{\text{sim}} ^2 + \rightarrow \log \rightarrow)$	$O(P_{\text{sim}} ^2 \rightarrow)$
Time-Céc [2]	$O(P_{\text{sim}} \Sigma \log \Sigma + \rightarrow \log \rightarrow)$	$O(P_{\text{sim}} \rightarrow)$
GPP [3]	$O(P_{\text{sim}} ^2 \log P_{\text{sim}} + \Sigma \log \Sigma)$	$O(P_{\text{sim}} ^2 \rightarrow)$
Mar [8]	$O((\Sigma + P_{\text{sim}} ^2) \log P_{\text{sim}})$	$O(\rightarrow + P_{\text{sim}} \Sigma + P_{\text{sim}} ^3)$
RT [12]	$O(P_{\text{sim}} \Sigma \log \Sigma)$	$O(P_{\text{sim}} \rightarrow)$
ESim (this paper)	$O(P_{\text{sim}} ^2 \log P_{\text{sim}} + \Sigma \log \Sigma)$	$O(P_{\text{sim}} \rightarrow \log \Sigma)$

We remark that all the above space bounds are bit space complexities, i.e., the word size is a single bit. Let us also remark that both articles [3,4] state that the bit space complexity of GPP is in $O(|P_{\text{sim}}|^2 + |\Sigma| \log |P_{\text{sim}}|)$. However, as observed also in [2], this is not precise. In fact, the algorithm GPP [3, Section 4, p. 98] assumes that the states belonging to some block are stored as a doubly linked list, and this entails a bit space complexity in $O(|\Sigma| \log |\Sigma|)$. Furthermore, GPP uses Henzinger, Henzinger and Kopke [5] simulation algorithm (HKK) as a subroutine, whose bit space complexity is in $O(|\Sigma|^2 \log |\Sigma|)$, which is called on a Kripke structure where states are blocks of the current partition. The bit space complexity of GPP must therefore include an additive term $|P_{\text{sim}}|^2 \log |P_{\text{sim}}|$ and therefore results to be $O(|P_{\text{sim}}|^2 \log |P_{\text{sim}}| + |\Sigma| \log |\Sigma|)$. It is worth observing that a space complexity in $O(|P_{\text{sim}}|^2 + |\Sigma| \log |P_{\text{sim}}|)$ can be considered optimal for a simulation algorithm, since this is of the same order as the size of the output, which needs $|P_{\text{sim}}|^2$ space for storing the simulation preorder as a partial order on simulation equivalence classes and $|\Sigma| \log |P_{\text{sim}}|$ space for storing the simulation equivalence class for any state. Hence, the bit space complexities of GPP and Space-Céc can be considered quasi-optimal. As far as time complexity is concerned, the algorithms RT and Time-Céc both feature the best time bound $O(|P_{\text{sim}}| \rightarrow)$.

We present here a novel space and time Efficient Simulation algorithm, called ESIm, which features a time complexity in $O(|P_{\text{sim}}| \rightarrow \log |\Sigma|)$ and a bit space complexity in $O(|P_{\text{sim}}|^2 \log |P_{\text{sim}}| + |\Sigma| \log |\Sigma|)$. Thus, ESIm reaches the best space bound of GPP and significantly improves the GPP time bound $O(|P_{\text{sim}}|^2 \rightarrow)$ by replacing a multiplicative factor $|P_{\text{sim}}|$ with $\log |\Sigma|$. Furthermore, ESIm significantly improves the RT space bound $O(|P_{\text{sim}}| |\Sigma| \log |\Sigma|)$ and closely approaches the best time bound $O(|P_{\text{sim}}| \rightarrow)$ of RT and Time-Céc.

ESIm is a partition refinement algorithm, meaning that it maintains and iteratively refines a so-called partition-relation pair $\langle P, \trianglelefteq \rangle$, where P is a partition of Σ that overapproximates the final simulation partition P_{sim} , while \trianglelefteq is a binary relation over P which overapproximates the final simulation preorder. ESIm relies on the following three main points, which in particular allow to attain the above complexity bounds.

- (1) Two distinct notions of partition and relation stability for a partition-relation pair are introduced. Accordingly, at a logical level, ESIm is designed as a partition refinement algorithm which iteratively performs two clearly distinct refinement steps: the refinement of the current partition P which splits some blocks of P and the refinement of the relation \trianglelefteq which removes some pairs of blocks from \trianglelefteq .
- (2) ESIm exploits a logical characterization of partition refiners, i.e. blocks of P that allow to split the current partition P , which admits an efficient implementation.
- (3) ESIm only relies on data structures, like lists and matrices, that are indexed on and contain blocks of the current partition P . The hard task here is to devise efficient ways to keep updated these partition-based data structures along the iterations of ESIm. We show that this can be done efficiently, in particular by resorting to Hopcroft's "process the smaller half" principle [7] when updating a crucial data structure after a partition split.

Due to lack of space, some auxiliary algorithms and the proofs of all the results are omitted.

2 Background

Notation. If $R \subseteq \Sigma \times \Sigma$ is any relation and $X \subseteq \Sigma$ then $R(X) \triangleq \{x' \in \Sigma \mid \exists x \in X. (x, x') \in R\}$. Recall that R is a preorder relation when it is reflexive and transitive. If f is a function defined on $\wp(\Sigma)$ and $x \in \Sigma$ then we often write $f(x)$ to mean $f(\{x\})$. $\text{Part}(\Sigma)$ denotes the set of partitions of Σ . If $P \in \text{Part}(\Sigma)$, $s \in \Sigma$ and $S \subseteq \Sigma$ then $P(s)$ denotes the block of P that contains s while $P(S) = \cup_{s \in S} P(s)$. $\text{Part}(\Sigma)$ is endowed with the standard partial order \preceq : $P_1 \preceq P_2$, i.e. P_2 is coarser than P_1 , iff for any $s \in \Sigma$, $P_1(s) \subseteq P_2(s)$. If $P_1 \preceq P_2$ and $B \in P_1$ then $P_2(B)$ is a block of P_2 which is also denoted by $\text{parent}_{P_2}(B)$. For a given nonempty subset $S \subseteq \Sigma$ called splitter, we denote by $\text{Split}(P, S)$ the partition obtained from P by replacing each block $B \in P$ with $B \cap S$ and $B \setminus S$, where we also allow no splitting, namely $\text{Split}(P, S) = P$ (this happens exactly when $P(S) = S$).

Simulation Preorder and Equivalence. A transition system (Σ, \rightarrow) consists of a set Σ of states and of a transition relation $\rightarrow \subseteq \Sigma \times \Sigma$. Given a set AP of atoms (of some specification language), a Kripke structure (KS) $\mathcal{K} = (\Sigma, \rightarrow, \ell)$ over AP consists of a transition system (Σ, \rightarrow) together with a state labeling function $\ell : \Sigma \rightarrow \wp(AP)$. The state partition induced by ℓ is denoted by $P_\ell \triangleq \{\{s' \in \Sigma \mid \ell(s) = \ell(s')\} \mid s \in \Sigma\}$. The predecessor/successor transformers $\text{pre}, \text{post} : \wp(\Sigma) \rightarrow \wp(\Sigma)$ are defined as usual: $\text{pre}(T) \triangleq \{s \in \Sigma \mid \exists t \in T. s \rightarrow t\}$ and $\text{post}(S) \triangleq \{t \in \Sigma \mid \exists s \in S. s \rightarrow t\}$. If $S_1, S_2 \subseteq \Sigma$ then $S_1 \rightarrow^{\exists} S_2$ iff there exist $s_1 \in S_1$ and $s_2 \in S_2$ such that $s_1 \rightarrow s_2$.

A relation $R \subseteq \Sigma \times \Sigma$ is a simulation on a Kripke structure $(\Sigma, \rightarrow, \ell)$ if for any $s, s' \in \Sigma$, if $s' \in R(s)$ then:

- (A) $\ell(s) = \ell(s')$;
- (B) for any $t \in \Sigma$ such that $s \rightarrow t$, there exists $t' \in \Sigma$ such that $s' \rightarrow t'$ and $t' \in R(t)$.

Given $s, t \in \Sigma$, t simulates s , denoted by $s \leq t$, if there exists a simulation relation R such that $t \in R(s)$. It turns out that the largest simulation on a given KS exists, is a preorder relation called simulation preorder and is denoted by R_{sim} . Thus, for any $s, t \in \Sigma$, $s \leq t$ iff $(s, t) \in R_{\text{sim}}$. Simulation equivalence $R_{\text{sim}eq}$ is the symmetric reduction of R_{sim} , namely $R_{\text{sim}eq} \triangleq R_{\text{sim}} \cap R_{\text{sim}}^{-1}$, so that $(s, t) \in R_{\text{sim}eq}$ iff $s \leq t$ and $t \leq s$. $P_{\text{sim}} \in \text{Part}(\Sigma)$ denotes the partition corresponding to the equivalence $R_{\text{sim}eq}$ and is called the simulation partition.

3 Logical Simulation Algorithm

A partition-relation pair $\mathcal{P} = \langle P, \trianglelefteq \rangle$, PR for short, is a state partition $P \in \text{Part}(\Sigma)$ together with a binary relation $\trianglelefteq \subseteq P \times P$ between blocks of P . We write $B \triangleleft C$ when $B \trianglelefteq C$ and $B \neq C$ and $(B', C') \trianglelefteq (B, C)$ when $B' \trianglelefteq B$ and $C' \trianglelefteq C$. When \trianglelefteq is a preorder/partial order then \mathcal{P} is called, respectively, a preorder/partial order PR.

PRs allow to represent symbolically, i.e. through state partitions, a relation between states. A relation $R \subseteq \Sigma \times \Sigma$ induces a PR $\text{PR}(R) = \langle P, \trianglelefteq \rangle$ defined as follows:

$$\forall s \in \Sigma. P(s) \triangleq \{t \in \Sigma \mid R(s) = R(t)\}; \quad \forall s, t \in \Sigma. P(s) \trianglelefteq P(t) \text{ iff } t \in R(s).$$

It is easy to note that if R is a preorder then $\text{PR}(R)$ is a partial order PR. On the other hand, a PR $\mathcal{P} = \langle P, \trianglelefteq \rangle$ induces the following relation $\text{Rel}(\mathcal{P}) \subseteq \Sigma \times \Sigma$:

$$(s, t) \in \text{Rel}(\mathcal{P}) \Leftrightarrow P(s) \trianglelefteq P(t).$$

Here, if \mathcal{P} is a preorder PR then $\text{Rel}(\mathcal{P})$ is clearly a preorder.

A PR $\mathcal{P} = \langle P, \trianglelefteq \rangle$ is defined to be a simulation PR on a KS \mathcal{K} when $\text{Rel}(\mathcal{P})$ is a simulation on \mathcal{K} , namely when \mathcal{P} represents a simulation relation between states. Hence, if \mathcal{P} is a simulation PR and $P(s) = P(t)$ then s and t are simulation equivalent, while if $P(s) \trianglelefteq P(t)$ then t simulates s .

Given a PR $\mathcal{P} = \langle P, \trianglelefteq \rangle$, the map $\mu_{\mathcal{P}} : \wp(\Sigma) \rightarrow \wp(\Sigma)$ is defined as follows:

$$\text{for any } X \in \wp(\Sigma), \mu_{\mathcal{P}}(X) \triangleq \text{Rel}(\mathcal{P})(X) = \cup\{C \in P \mid \exists s \in X. P(s) \trianglelefteq C\}.$$

Note that, for any $s \in \Sigma$, $\mu_{\mathcal{P}}(s) = \mu_{\mathcal{P}}(P(s)) = \cup\{C \in P \mid P(s) \trianglelefteq C\}$. For preorder PRs, this map allows us to characterize the property of being a simulation PR as follows.

Theorem 3.1. *Let $\mathcal{P} = \langle P, \trianglelefteq \rangle$ be a preorder PR. Then, \mathcal{P} is a simulation iff*

- (i) *if $B \trianglelefteq C$, $b \in B$ and $c \in C$ then $\ell(b) = \ell(c)$;*
- (ii) *if $B \rightarrow^{\exists} C$ and $B \trianglelefteq D$ then $D \rightarrow^{\exists} \mu_{\mathcal{P}}(C)$;*
- (iii) *for any $C \in P$, $P = \text{Split}(P, \text{pre}(\mu_{\mathcal{P}}(C)))$.*

By Theorem 3.1, assuming that condition (i) holds, there are two possible reasons for a PR $\mathcal{P} = \langle P, \trianglelefteq \rangle$ for not being a simulation:

- (1) There exist $B, C, D \in P$ such that $B \rightarrow^{\exists} C$, $B \trianglelefteq D$, but $D \not\rightarrow^{\exists} \mu_{\mathcal{P}}(C)$; in this case we say that the block C is a *relation refiner* for \mathcal{P} .
- (2) There exist $B, C \in P$ such that $B \cap \text{pre}(\mu_{\mathcal{P}}(C)) \neq \emptyset$ and $B \setminus \text{pre}(\mu_{\mathcal{P}}(C)) \neq \emptyset$; in this case we say that the block C is a *partition refiner* for \mathcal{P} .

We therefore define $\text{RRRefiner}(\mathcal{P})$ and $\text{PRefiner}(\mathcal{P})$ as the sets of blocks of P that are, respectively, relation and partition refiners for \mathcal{P} . Accordingly, \mathcal{P} is defined to be relation or partition *stable* when, respectively, $\text{RRRefiner}(\mathcal{P}) = \emptyset$ or $\text{PRefiner}(\mathcal{P}) = \emptyset$. Then, Theorem 3.1 can be read as follows: \mathcal{P} is a simulation iff \mathcal{P} satisfies condition (i) and is both relation and partition stable.

If $C \in \text{PRefiner}(\mathcal{P})$ then P is first refined to $P' \triangleq \text{Split}(P, \text{pre}(\mu_{\mathcal{P}}(C)))$, i.e. P is split w.r.t. the splitter $S = \text{pre}(\mu_{\mathcal{P}}(C))$. Accordingly, the relation \trianglelefteq on P is transformed into the following relation \trianglelefteq' defined on P' :

$$\trianglelefteq' \triangleq \{(D, E) \in P' \times P' \mid \text{parent}_P(D) \trianglelefteq \text{parent}_P(E)\} \quad (\dagger)$$

Hence, two blocks D and E of the refined partition P' are related by \trianglelefteq' if their parent blocks $\text{parent}_P(D)$ and $\text{parent}_P(E)$ in P were related by \trianglelefteq . Hence, if $\mathcal{P}' = \langle P', \trianglelefteq' \rangle$ then for all $D \in P'$, we have that $\mu_{\mathcal{P}'}(D) = \mu_{\mathcal{P}}(\text{parent}_P(D))$. We will show that this refinement of $\langle P, \trianglelefteq \rangle$ is correct because if $B \in P$ is split into $B \setminus S$ and $B \cap S$ then all the states in $B \setminus S$ are not simulation equivalent to all the states in $B \cap S$. Note that if $B \in P$ has been split into $B \cap S$ and $B \setminus S$ then both $B \cap S \trianglelefteq' B \setminus S$ and $B \setminus S \trianglelefteq' B \cap S$ hold, and consequently \mathcal{P}' becomes relation unstable.

```

1 ESim(PR  $\langle P, \sqsubseteq \rangle$ )
2   Initialize(); PStabilize(); bool PStable := RStabilize(); bool RStable := tt;
3   while  $\neg(\text{PStable} \ \& \ \text{RStable})$  do
4     if  $\neg \text{PStable}$  then {RStable := PStabilize(); PStable := tt;}
5     if  $\neg \text{RStable}$  then {PStable := RStabilize(); RStable := tt;}

6 bool PStabilize()
7   Pold := P;
8   while  $\exists C \in \text{PRefiner}(\mathcal{P})$  do
9      $S := \text{pre}(\mu_{\mathcal{P}}(C)); P := \text{Split}(S);$ 
10    forall  $(D, E) \in P \times P$  do  $D \sqsubseteq E := \text{parent}_P(D) \sqsubseteq \text{parent}_P(E);$ 
11  return  $(P = P_{\text{old}});$ 

12 bool RStabilize()
13  // Precondition: PStable = tt
14   $\sqsubseteq_{\text{old}} := \sqsubseteq; \text{Delete} := \emptyset;$ 
15  while  $\exists C \in \text{RRefiner}(\mathcal{P})$  do
16     $\text{Delete} := \text{Delete} \cup \{(B, D) \in P \times P \mid B \sqsubseteq D, B \rightarrow^{\exists} C, D \not\rightarrow^{\exists} \mu_{\mathcal{P}}(C)\};$ 
17   $\sqsubseteq := \sqsubseteq \setminus \text{Delete};$ 
18  return  $(\sqsubseteq = \sqsubseteq_{\text{old}});$ 

```

Fig. 1. Logical Simulation Algorithm.

On the other hand, if \mathcal{P} is partition stable and $C \in \text{RRefiner}(\mathcal{P})$ then we will show that \sqsubseteq can be safely refined to the following relation \sqsubseteq' :

$$\begin{aligned} \sqsubseteq' &\triangleq \sqsubseteq \setminus \{(B, D) \in P \times P \mid B \rightarrow^{\exists} C, B \sqsubseteq D, D \not\rightarrow^{\exists} \mu_{\mathcal{P}}(C)\} \\ &= \{(B, D) \in P \times P \mid B \sqsubseteq D, (B \rightarrow^{\exists} C \Rightarrow D \rightarrow^{\exists} \mu_{\mathcal{P}}(C))\} \end{aligned} \quad (\ddagger)$$

because if $(B, D) \in \sqsubseteq \setminus \sqsubseteq'$ then all the states in D cannot simulate all the states in B .

The above facts lead us to design a basic simulation algorithm ESim described in Figure 1. ESim maintains a PR $\mathcal{P} = \langle P, \sqsubseteq \rangle$, which initially is $\langle P_{\ell}, \text{id} \rangle$ and is iteratively refined as follows:

PStabilize(): If $\langle P, \sqsubseteq \rangle$ is not partition stable then the partition P is split for $\text{pre}(\mu_{\mathcal{P}}(C))$ as long as a partition refiner C for \mathcal{P} exists, and when this happens the relation \sqsubseteq is transformed to \sqsubseteq' as defined by (\ddagger) ; at the end of this process, we obtain a PR $\mathcal{P}' = \langle P', \sqsubseteq' \rangle$ which is partition stable and if P has been actually refined, i.e. $P' \prec P$ then the current PR \mathcal{P}' becomes relation unstable.

RStabilize(): If $\langle P, \sqsubseteq \rangle$ is not relation stable then the relation \sqsubseteq is refined to \sqsubseteq' as described by (\ddagger) as long as a relation refiner for \mathcal{P} exists; hence, at the end of this refinement process $\langle P, \sqsubseteq' \rangle$ becomes relation stable but possibly partition unstable.

Moreover, the following properties of the current PR of ESim hold.

Lemma 3.2. *In any run of ESim, the following two conditions hold:*

- (i) If $PStabilize()$ is called on a partial order PR $\langle P, \trianglelefteq \rangle$ then at the exit we obtain a PR $\langle P', \trianglelefteq' \rangle$ which is a preorder.
- (ii) If $RStabilize()$ is called on a preorder PR $\langle P, \trianglelefteq \rangle$ then at the exit we obtain a PR $\langle P, \trianglelefteq' \rangle$ which is a partial order.

The main loop of ESIm terminates when the current PR $\langle P, \trianglelefteq \rangle$ becomes both partition and relation stable. By the above Lemma 3.2, the output PR \mathcal{P} of ESIm is a partial order, and hence a preorder, so that Theorem 3.1 can be applied to \mathcal{P} which then results to be a simulation PR. It turns out that this algorithm is correct, meaning that the output PR \mathcal{P} actually represents the simulation preorder.

Theorem 3.3 (Correctness). *Let Σ be finite. ESIm is correct, i.e., ESIm terminates on any input and if $\langle P, \trianglelefteq \rangle$ is the output PR of ESIm on input $\langle P_\ell, \text{id} \rangle$ then for any $s, t \in \Sigma$, $s \leq t \Leftrightarrow P(s) \trianglelefteq P(t)$.*

4 Efficient Implementation

4.1 Data Structures

ESIm is implemented by relying on the following data structures.

States: A state s is represented by a record that contains the list $\text{post}(s)$ of its successors, a pointer $s.\text{block}$ to the block $P(s)$ that contains s and a boolean flag used for marking purposes. The whole state space Σ is represented as a doubly linked list of states. $\{\text{post}(s)\}_{s \in \Sigma}$ therefore represents the input transition system.

Partition: The states of any block B of the current partition P are consecutive in the list Σ , so that B is represented by two pointers begin and end : $B.\text{begin}$ is the first state of B in Σ and $B.\text{end}$ is the successor of the last state of B in Σ , i.e., $B = [B.\text{begin}, B.\text{end}[$. Moreover, B stores a boolean flag $B.\text{intersection}$ and a block pointer $B.\text{brother}$ whose meanings are as follows: after a call to $Split(P, S)$ for splitting P w.r.t. a set of states S , if $B_1 = B \cap S$ and $B_2 = B \setminus S$, for some $B \in P$ that has been split by S then $B_1.\text{intersection} = \mathbf{tt}$ and $B_2.\text{intersection} = \mathbf{ff}$, while $B_1.\text{brother}$ points to B_2 and $B_2.\text{brother}$ points to B_1 . If instead B has not been split by S then $B.\text{intersection} = \mathbf{null}$ and $B.\text{brother} = \mathbf{null}$. Also, any block B stores in $Rem(B)$ a list of blocks of P , which is used by $RStabilize()$, and in $B.\text{preE}$ the list of blocks $C \in P$ such that $C \rightarrow^{\exists} B$. Finally, any block B stores in $B.\text{size}$ the size of B , in $B.\text{count}$ an integer counter bounded by $|P|$ which is used by $PStabilize()$ and a pair of boolean flags used for marking purposes. The current partition P is stored as a doubly linked list of blocks.

Relation: The current relation \trianglelefteq on P is stored as a resizable $|P| \times |P|$ boolean matrix. Recall that insert operations in a resizable array (whose capacity is doubled as needed) take amortized constant time and that a resizable matrix (or table) can be implemented as a resizable array of resizable arrays. The boolean matrix \trianglelefteq is resized by adding a new entry to \trianglelefteq , namely a new row and a new column, for any block B that is split into two new blocks $B \setminus S$ and $B \cap S$. The old entry B becomes the entry for the new block $B \setminus S$ while the new entry is used for the new block $B \cap S$.

```

1 bool PStabilize()
2 list(Block) split := ∅;
3 while (C := FindPRefiner()) ≠ null do
4   list(State) S := preμ(C); split := Split(S); updateRel(split);
5   updateBCount(split); updatePreE(); updateCount(split); updateRem(split);
6 return (split = ∅);

7 Block FindPRefiner()
8 forall B ∈ P do
9   list(Block) p := Post(B);
10  forall C ∈ p do if (Count(B, C) = 1) then return C;
11 return null;

12 list(Block) Post(Block B)
13 list(Block) p := ∅;
14 forall b ∈ B, do
15   forall c ∈ post(b) do
16     Block C := c.block;
17     if unmarked1(C) then {mark1(C); C.count = 0; p.append(C);}
18     if unmarked2(C) then {mark2(C); C.count++;}
19   forall C ∈ p do unmark2(C);
20 forall C ∈ p do {unmark1(C); if (C.count = B.size) then p.remove(C);}
21 return p;

```

Fig. 2. PStabilize() Algorithm.

Auxiliary Data Structures: We store and maintain a resizable boolean matrix BCount and a resizable integer matrix Count, both indexed over P , whose meanings are as follows:

$$\text{BCount}(B, C) \triangleq \begin{cases} 1 & \text{if } B \rightarrow^{\exists} C \\ 0 & \text{if } B \not\rightarrow^{\exists} C \end{cases} \quad \text{Count}(B, C) \triangleq \sum_{E \triangleright C} \text{BCount}(B, E).$$

Hence, $\text{Count}(B, C)$ stores the number of blocks E such that $C \trianglelefteq E$ and $B \rightarrow^{\exists} E$. The table Count allows to implement the test $B \not\rightarrow^{\exists} \text{pre}(\mu_P(C))$ in constant time as $\text{Count}(B, C) = 0$.

The data structures BCount, preE, Count and Rem are initialized by a function Initialize() at line 2 of ESIm, which is here omitted.

4.2 Partition Stability

Our implementation of ESIm will exploit the following logical characterization of partition refiners.

Theorem 4.1. *Let $\langle P, \trianglelefteq \rangle$ be a partial order PR. Then, $\text{PRefiner}(\langle P, \trianglelefteq \rangle) \neq \emptyset$ iff there exist $B, C \in P$ such that the following three conditions hold:*

- (i) $B \rightarrow^{\exists} C$;
- (ii) for any $C' \in P$, if $C \triangleleft C'$ then $B \not\rightarrow^{\exists} C'$;
- (iii) $B \not\subseteq \text{pre}(C)$.

Notice that this characterization of partition refiners requires that the current PR is a partial order relation and, by Lemma 3.2, for any call to $PStabilize()$, this is actually guaranteed by the ESim algorithm.

The algorithm in Figure 2 is an implementation of the $PStabilize()$ function that relies on Theorem 4.1 and on the above data structures. The function $FindPRefiner()$ implements the conditions of Theorem 4.1: it returns a partition refiner for the current PR $\mathcal{P} = \langle P, \preceq \rangle$ when this exists, otherwise it returns a null pointer. Given a block $B \in P$, the function $Post(B)$ returns a list of blocks $C \in P$ that satisfy conditions (i) and (iii) of Theorem 4.1, i.e., those blocks C such that $B \rightarrow^{\exists} C$ and $B \not\subseteq \text{pre}(C)$. This is accomplished through the counter $C.\text{count}$ that at the exit of the for-loop at lines 14-19 in Figure 2 stores the number of states in B having (at least) an outgoing transition to C , i.e., $C.\text{count} = |B \cap \text{pre}(C)|$. Hence, we have that:

$$B \rightarrow^{\exists} C \text{ and } B \not\subseteq \text{pre}(C) \Leftrightarrow 1 \leq C.\text{count} < B.\text{size}.$$

Then, for any candidate partition refiner $C \in Post(B)$, it remains to check condition (ii) of Theorem 4.1. This condition is checked in $FindPRefiner()$ by testing whether $\text{Count}(B, C) = 1$: this is correct because $\text{Count}(B, C) \geq 1$ holds since $C \in Post(B)$ and therefore $B \rightarrow^{\exists} C$, so that

$$\text{Count}(B, C) = 1 \text{ iff } \forall C' \in P. C \triangleleft C' \Rightarrow B \not\rightarrow^{\exists} C'.$$

Hence, if $\text{Count}(B, C) = 1$ holds at line 10 of $FindPRefiner()$, by Theorem 4.1, C is a partition refiner. Once a partition refiner C has been returned by $Post(B)$, $PStabilize()$ splits the current partition P w.r.t. the splitter $S = \text{pre}(\mu_{\mathcal{P}}(C))$ by calling the function $Split(S)$, updates the relation \preceq as defined by equation (†) in Section 3 by calling $updateRel()$, updates the data structures BCount, preE, Count and Rem, and then check again whether a partition refiner exists. At the exit of the main while-loop of $PStabilize()$, the current PR $\langle P, \preceq \rangle$ is partition stable.

$PStabilize()$ calls the functions $\text{pre}\mu()$ and $Split()$ that are here omitted. Recall that the states of a block B of P are consecutive in the list of states Σ , so that B is represented as $B = [B.\text{begin}, B.\text{end}]$. The implementation of $Split(S)$ is quite standard (see e.g. [12]): this is based on a linear scan of the states in S and for each state in S performs some constant time operations. Hence, $Split(S)$ takes $O(|S|)$ time. Also, $Split(S)$ returns the list $split$ of blocks $B \setminus S$ such that $\emptyset \subsetneq B \setminus S \subsetneq B$ (i.e., $B.\text{intersection} = \mathbf{ff}$). Let us remark that a call $Split(S)$ may affect the ordering of the states in the list Σ because states are moved from old blocks to newly generated blocks.

We will show that the overall time complexity of $PStabilize()$ along a whole run of ESim is in $O(|P_{\text{sim}}| \rightarrow | \cdot |)$.

4.3 Updating Data Structures

In the function $PStabilize()$, after calling $Split(S)$, firstly we need to update the boolean matrix that stores the relation \preceq in accordance with definition (†) in Section 3. After

that, since both P and \sqsubseteq are changed we need to update the data structures BCount , preE , Count and Rem . We omit the implementations of the functions $\text{updateRel}()$, $\text{updateBCount}()$, $\text{updatePreE}()$ and $\text{updateRem}()$, which are quite straightforward.

The function $\text{updateCount}()$ is in Figure 3 and deserves special care in order to design a time efficient implementation. The core of the $\text{updateCount}()$ algorithm follows Hopcroft’s “process the smaller half” principle [7] for updating the integer matrix Count . Let P' be the partition which is obtained by splitting the partition P w.r.t. the splitter S . Let B be a block of P that has been split into $B \cap S$ and $B \setminus S$. Thus, we need to update $\text{Count}(B \cap S, C)$ and $\text{Count}(B \setminus S, C)$ for any $C \in P'$ by knowing $\text{Count}(B, \text{parent}_P(C))$. Let us first observe that after lines 4-6 of $\text{updateCount}()$, we have that for any $B, C \in P'$, $\text{Count}(B, C) = \text{Count}(\text{parent}_P(B), \text{parent}_P(C))$. Let X be the block in $\{B \cap S, B \setminus S\}$ with the smaller size, and let Z be the other block, so that $|X| \leq |B|/2$ and $|X| + |Z| = |B|$. Let C be any block in P' . We set $\text{Count}(X, C)$ to 0, while $\text{Count}(Z, C)$ is left unchanged, namely $\text{Count}(Z, C) = \text{Count}(B, C)$. We can correctly update both $\text{Count}(Z, C)$ and $\text{Count}(X, C)$ by just scanning all the outgoing transitions from X . In fact, if $x \in X$, $x \rightarrow y$ and the block $P(y)$ is scanned for the first time then for all $C \sqsubseteq P(y)$, $\text{Count}(X, C)$ is incremented by 1, while if $Z \not\rightarrow^{\exists} P(y)$, i.e. $\text{BCount}(Z, P(y)) = 0$, then $\text{Count}(Z, C)$ is decremented by 1. The correctness of this procedure goes as follows:

- (1) At the end, $\text{Count}(X, C)$ is clearly correct because its value has been re-computed from scratch.
- (2) At the end, $\text{Count}(Z, C)$ is correct because $\text{Count}(Z, C)$ initially stores the value $\text{Count}(B, C)$, and if there exists some block D such that $C \sqsubseteq D$, $B \rightarrow^{\exists} D$ whereas $Z \not\rightarrow^{\exists} D$ — this is correctly implemented at line 19 as $\text{BCount}(Z, D) = 0$, since the data structure BCount is up to date — then necessarily $X \rightarrow^{\exists} D$, because B has been split into X and Z , so that $D = P(y)$ for some $y \in \text{post}(X)$, namely D has been taken into account by some increment $\text{Count}(X, C)++$ and consequently $\text{Count}(Z, C)$ is decremented by 1 at line 19.

Moreover, if some block $D \in P' \setminus \{B \cap S, B \setminus S\}$ is such that both $D \rightarrow^{\exists} X$ and $D \rightarrow^{\exists} Z$ hold then for all the blocks $C \in P$ such that $C \sqsubseteq X$ (or, equivalently, $C \sqsubseteq Z$), we need to increment $\text{Count}(D, C)$ by 1. This is done at lines 21-22 by relying on the updated data structures preE and BCount .

Let us observe that the time complexity of a single call of $\text{updateCount}(\text{split})$ is

$$|P|(|\text{split}| + \sum_{X \in \text{split}} (|\{(x, y) \mid x \in X, y \in \Sigma, x \rightarrow y\}| + |\{(X, D) \mid D \in P, X \rightarrow^{\exists} D\}|)).$$

Hence, let us calculate the overall time complexity of $\text{updateCount}()$. If X and X' are two blocks that are scanned in two different calls of updateCount and $X' \subseteq X$ then $|X'| \leq |X|/2$. Consequently, any transition $x \rightarrow y$ at line 16 and $D \rightarrow^{\exists} X$ at line 21 can be scanned in some call of $\text{updateCount}()$ at most $\log_2 |\Sigma|$ times. Thus, the overall time complexity of $\text{updateCount}()$ is in $O(|P_{\text{sim}}| \rightarrow |\log |\Sigma|)$.

4.4 Relation Stability

The basic procedure $\text{RStabilize}()$ in Figure 1 is implemented by the algorithm in Figure 4. Let $\mathcal{P}^{\text{in}} = \langle P, \sqsubseteq^{\text{in}} \rangle$ be the current PR when calling $\text{RStabilize}()$. For each rela-

```

1 // Precondition: BCount and preE are updated with the current PR
2 updateCount(list<Block> split)
3   forall B ∈ split do addNewEntry(B) in matrix Count;
4   forall B ∈ P, C ∈ split do
5     if (B.intersection = tt) then Count(B, C) := Count(B.brother, C.brother);
6     else Count(B, C) := Count(B, C.brother);
7   forall C ∈ P, B ∈ split do
8     if (C.intersection = ff) then Count(B, C) := Count(B.brother, C);
9   forall C ∈ P do unmark(C);
10  forall B ∈ split do
11    // Update Count(B, ·) and Count(B.brother, ·)
12    Block X, Z;
13    if (B.size ≤ B.brother.size) then {X := B; Z := B.brother;}
14    else {X := B.brother; Z := B;}
15    forall C ∈ P do {Count(X, C) := 0; /* Count(Z, C) := Count(B, C); */}
16    forall x ∈ X, y ∈ post(x) such that unmarked(y.block) do
17      mark(y.block);
18      forall C ∈ P such that C ⊆ y.block do
19        Count(X, C)++; if (BCount(Z, y.block) = 0) then Count(Z, C)--;
20    // For all D ∉ {B, B.brother}, updateCount(D, ·)
21    forall D ∈ X.preE such that (D ≠ X & D ≠ Z & BCount(D, Z) = 1) do
22      forall C ∈ P such that C ⊆ X do Count(D, C)++;

```

Fig. 3. *updateCount()* function.

tion refiner $C \in P$, $RStabilize()$ must iteratively refine the initial relation $\trianglelefteq^{\text{in}}$ in accordance with equation (‡) in Section 3. Hence, if $B \rightarrow^{\exists} C$, $B \trianglelefteq D$ and $D \not\stackrel{\exists}{\mu}_{\mathcal{P}^{\text{in}}}(C)$, the entry $B \trianglelefteq D$ of the boolean matrix that represents the relation \trianglelefteq must be set to **ff**. Thus, the idea is to store and incrementally maintain for each block $C \in P$ a list $Rem(C)$ of blocks $D \in P$ such that: (A) If C is a relation refiner for \mathcal{P}^{in} then $Rem(C) \neq \emptyset$; (B) If $D \in Rem(C)$ then necessarily $D \not\stackrel{\exists}{\mu}_{\mathcal{P}^{\text{in}}}(C)$. It turns out that C is a relation refiner for \mathcal{P}^{in} iff there exist blocks B and D such that $B \rightarrow^{\exists} C$, $D \in Rem(C)$ and $B \trianglelefteq D$. Hence, the set of blocks $Rem(C)$ is reminiscent of the set of states $remove(s)$ used in Henzinger et al.'s [5] simulation algorithm, since each pair (B, D) which must be removed from the relation \trianglelefteq is such that $D \in Rem(C)$, for some block C .

Initially, namely at the first call of $RStabilize()$ by $ESim$, $Rem(C)$ is set by the function $Initialize()$ to $\{D \in P \mid D \rightarrow^{\exists} \Sigma, D \not\stackrel{\exists}{\mu}_{\mathcal{P}}(C)\}$. Hence, $RStabilize()$ scans all the blocks in the current partition P and selects those blocks C such that $Rem(C) \neq \emptyset$, which are therefore candidate to be relation refiners. Then, by scanning all the blocks $B \in C.preE$ and $D \in Rem(C)$, if $B \trianglelefteq D$ holds then the entry $B \trianglelefteq D$ must be set to **ff**. However, the removal of the pair (B, D) from the current relation \trianglelefteq may affect the function $\mu_{\mathcal{P}}$. This is avoided by making a copy $oldRem(C)$ of all the $Rem(C)$'s at the beginning of $RStabilize()$ and then using this copy. During the main for-loop of

```

1 bool RStabilize()
2 //  $\mu_{\mathcal{P}}^{\text{in}} := \mu_{\mathcal{P}}$ ;
3 forall  $C \in P$  do {oldRem( $C$ ) := Rem( $C$ ); Rem( $C$ ) =  $\emptyset$ ; }
4 bool Removed := ff;
5 forall  $C \in P$  such that oldRem( $C$ )  $\neq \emptyset$  do
6     // Invariant (Inv):  $\forall C \in P. \text{Rem}(C) = \{D \in P \mid D \rightarrow^{\exists} \mu_{\mathcal{P}}^{\text{in}}(C), D \not\rightarrow^{\exists} \mu_{\mathcal{P}}(C)\}$ 
7     forall  $B \in C.\text{preE}, D \in \text{oldRem}(C)$  such that ( $B \preceq D$ ) do
8          $B \preceq D$  := ff; Removed := tt;
9         // update Count and Rem
10        forall  $F \in D.\text{preE}$  do
11            Count( $F, B$ ) := Count( $F, B$ ) - 1;
12            if (Count( $F, B$ ) = 0 & Rem( $B$ ) =  $\emptyset$ ) then Rem( $B$ ).append( $F$ );
13 return  $\neg$ Removed;

```

Fig. 4. *RStabilize*() Algorithm.

RStabilize(), *Rem*(C) must satisfy the following invariant property:

$$(\text{Inv}): \forall C \in P. \text{Rem}(C) = \{D \in P \mid D \rightarrow^{\exists} \mu_{\mathcal{P}}^{\text{in}}(C), D \not\rightarrow^{\exists} \mu_{\mathcal{P}}(C)\}.$$

This means that at the beginning of *RStabilize*(), any *Rem*(C) is set to empty, and after the removal of a pair (B, D) from \preceq , since $\mu_{\mathcal{P}}(B)$ has changed, we need: (i) to update the matrix *Count*, for all the entries (F, B) where $F \rightarrow^{\exists} D$, and (ii) to check if there is some block F such that $F \not\rightarrow^{\exists} \mu_{\mathcal{P}}(B)$, because any such F must be added to *Rem*(B) in order to maintain the invariant property (Inv).

4.5 Complexity

The time complexity of the algorithm *ESim* relies on the following key properties:

- (1) The overall number of partition refiners found by *ESim* is in $O(|P_{\text{sim}}|)$. Moreover, the overall number of newly generated blocks by the splitting operations performed by calling *Split*(S) at line 4 of *PStabilize*() is in $O(|P_{\text{sim}}|)$. In fact, let $\{P_i\}_{i \in [0, n]}$ be the sequence of different partitions computed by *ESim* where P_0 is the initial partition P_{ℓ} , P_n is the final partition P_{sim} and for all $i \in [1, n]$, P_i is the partition after the i -th call to *Split*(S), so that $P_i \prec P_{i-1}$. The number of new blocks which are produced by a call *Split*(S) that refines P_i to P_{i+1} is $2(|P_{i+1}| - |P_i|)$. Thus, the overall number of newly generated blocks is $\sum_{i=1}^n 2(|P_i| - |P_{i-1}|) = 2(|P_{\text{sim}}| - |P_{\ell}|) \in O(|P_{\text{sim}}|)$.
- (2) The invariant (Inv) of the sets *Rem*(C) guarantees the following property: if C_1 and C_2 are two blocks that are selected by the for-loop at line 5 of *RStabilize*() in two different calls of *RStabilize*(), and $C_2 \subseteq C_1$ (possibly $C_1 = C_2$) then $(\cup \text{Rem}(C_1)) \cap (\cup \text{Rem}(C_2)) = \emptyset$.

Theorem 4.2 (Complexity). *ESim* runs in $O(|P_{\text{sim}}|^2 \log |P_{\text{sim}}| + |\Sigma| \log |\Sigma|)$ -space and $O(|P_{\text{sim}}| \rightarrow \log |\Sigma|)$ -time.

5 Further Work

We see a couple of interesting avenues for further work. A first natural question arises: can the time complexity of ESim be further improved and reaches the time complexity of RT? This would require to eliminate the multiplicative factor $\log |\Sigma|$ from the time complexity of ESim and, presently, this seems to us quite hard to achieve. More in general, it would be interesting to investigate whether some lower space and time bounds can be stated for the simulation preorder problem. Secondly, ESim is designed for Kripke structures. While an adaptation of a simulation algorithm from Kripke structures to labeled transition systems (LTSs) can be conceptually simple, unfortunately such a shift may lead to some loss in both space and time complexities, as argued in [2]. We mention the works [1,6] and [8] that provide simulation algorithms for LTSs by adapting, respectively, RT and GPP. It is thus worth investigating whether and how ESim can be efficiently adapted to work with LTSs.

Acknowledgements. We acknowledge the contribution of Francesco Tapparo to a preliminary stage of this research which was informally presented in [11]. This work was partially supported by Microsoft Research SEIF 2013 Award and by the University of Padova under the project BECOM.

References

1. P.A. Abdulla, A. Bouajjani, L. Holík, L. Kaati and T. Vojnar. Computing simulations over tree automata. In *Proc. 14th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'08)*, LNCS 4963, pp. 93–108, 2008.
2. G. Cécé. Three simulation algorithms for labelled transition systems. Preprint cs.arXiv:1301.1638, arXiv.org, 2013.
3. R. Gentilini, C. Piazza and A. Policriti. From bisimulation to simulation: coarsest partition problems. *J. Automated Reasoning*, 31(1):73-103, 2003.
4. R. van Glabbeek and B. Ploeger. Correcting a space-efficient simulation algorithm. In *Proc. 20th Int. Conf. on Computer Aided Verification (CAV'08)*, LNCS 5123, pp. 517-529, 2008.
5. M.R. Henzinger, T.A. Henzinger and P.W. Kopke. Computing simulations on finite and infinite graphs. In *Proc. 36th IEEE FOCS*, 453-462, 1995.
6. L. Holík and J. Šimáček. Optimizing an LTS-simulation algorithm. *Computing and Informatics*, 29(6+):1337-1348, 2010.
7. J.E. Hopcroft. A $n \log n$ algorithm for minimizing states in a finite automaton. In Z. Kohavi and A. Paz eds., *Theory of Machines and Computations*, pp. 189-176, Academic Press, 1971.
8. J. Markovski. Saving time in a space-efficient simulation algorithm. In *Proc. 11th Int. Conf. on Quality Software*, pp. 244-251, IEEE, 2011.
9. R. Paige and R.E. Tarjan. Three partition refinement algorithms. *SIAM J. Comput.*, 16(6):973-989, 1987
10. F. Ranzato and F. Tapparo. A new efficient simulation equivalence algorithm. In *Proc. 22nd Annual IEEE Symp. on Logic in Computer Science (LICS'07)*, pp. 171-180, 2007.
11. F. Ranzato and F. Tapparo. A time and space efficient simulation algorithm. Short talk at *24th Annual IEEE Symposium on Logic in Computer Science (LICS'09)*, 2009.
12. F. Ranzato and F. Tapparo. An efficient simulation algorithm based on abstract interpretation. *Inf. Comput.*, 208(1):1-22, 2010.