

An algorithm for Bisimilarity on finite-state LTSs

- Establishing bisimilarity is a special case of a more general problem, called *Relational Coarsest Partition Problem* (RCPP).
- RCPP is based on refining a set into classes (or blocks) by starting from the coarsest possible partition (i.e., the set itself), by separating elements that are not equivalent (according to a specific rule)
 - Start with the coarsest possible partition (i.e., the given set);
 - Split the blocks of the partition (by using the splitting rule) until possible;
 - OUTPUT: two states are in relation if and only if at the end they belong to the same block.
- The key feature is the rule that decides the splitting of blocks in sub-blocks
 - This characterizes the different equivalences that can be obtained by using the algorithm for RCPP.

- Let B be a generic block of a partition and s and t be two states belonging to B
- Let us call $block(s)$ the set of states belonging to the same block as s.
- According to bisimilarity, s and t must be separated (and, hence, B must be split) iff:

$$\exists \alpha, s' \text{ s.t. } s \xrightarrow{\alpha} s' \wedge \forall t' (t \not\xrightarrow{\alpha} t' \vee t' \notin block(s'))$$

- We shall use an alternative condition (equivalent to the splitting rule given above).
- Define $tgt_{\alpha}(s) = \{B' : \exists s'. s \xrightarrow{\alpha} s' \wedge s' \in B'\}$
- The alternative splitting condition is:

$$\exists \alpha . tgt_{\alpha}(s) \neq tgt_{\alpha}(t)$$

Main Algorithm

INPUT: $LTS=(S,\Lambda,\rightarrow)$, $s,t\in S$

```
    P ← {S}
LOOP: forall B ∈ P do
    for all α ∈ Λ do
        I ← split(B, α, P)
        if I ≠ {B} then
            P ← (P \ {B}) ∪ I
            goto LOOP
forall B ∈ P do
    if {s,t} ⊆ B then return true
return false
```

Splitting Algorithm

INPUT: (B, α, P)

$s \in B$

$B1 \leftarrow \emptyset \quad B2 \leftarrow \emptyset$

forall $t \in B$ do

 if $\text{tgt}_\alpha(s) = \text{tgt}_\alpha(t)$ then $B1 \leftarrow B1 \cup \{t\}$


 else $B2 \leftarrow B2 \cup \{t\}$

if $B2 = \emptyset$ then return $\{B\}$

 else return $\{B1, B2\}$


An Example

Assigned Value	State
1	$a \bar{a}$
2	$a\bar{a} + \bar{a}a + \tau$
3	\bar{a}
4	a
5	0



- $P = \{\{1,2,3,4,5\}\}$:
 - Choose a and $B = \{1,2,3,4,5\}$
 - $\text{tgt}_a(1) = \text{tgt}_a(2) = \text{tgt}_a(4) = \{B\}$; $\text{tgt}_a(3) = \text{tgt}_a(5) = \emptyset$
 - $\text{split}(B, a, P)$ replaces B by two sub-blocks in P
- $P = \{\{1,2,4\}, \{3,5\}\}$:
 - Choose \underline{a} and $B = \{1,2,4\}$
 - $\text{tgt}_{\underline{a}}(1) = \text{tgt}_{\underline{a}}(2) = \{B\}$, whereas $\text{tgt}_{\underline{a}}(4) = \emptyset$
 - $\text{split}(B, \underline{a}, P)$ replaces B by two sub-blocks in P
- $P = \{\{1,2\}, \{4\}, \{3,5\}\}$:
 - Choose \underline{a} and $B = \{3,5\}$
 - $\text{tgt}_{\underline{a}}(3) = \{B\}$ whereas $\text{tgt}_{\underline{a}}(5) = \emptyset$
 - $\text{split}(B, \underline{a}, P)$ replaces B by two sub-blocks in P
- With $P = \{\{1, 2\}, \{4\}, \{3\}, \{5\}\}$, every block and action lead to no other split
- Because 1 and 2 are in the same block, $a|\underline{a}$ and $a\underline{a} + \underline{a}a + \tau$ are bisimilar

- The algorithm can also be used to prove that two processes are NOT equivalent, e.g.



- Like before, we arrive at $P = \{\{1, 2\}, \{4\}, \{3\}, \{5\}\}$
- But now, if we choose $B = \{1, 2\}$ and action τ , we can further split B into $\{1\}$ and $\{2\}$
- So, $a|\bar{a}$ and $a\bar{a} + \bar{a}a$ are NOT bisimilar

Correctness Proof

- The algorithm always terminates
 - every cycle works on a finite number of elements
 - the cardinality of P is at most $|S|$ (one state in every block)
 - in this case, no block can be further split
 - this ensures termination!
- Let P_f be the partition at the end of the algorithm; then,
$$s \sim t \text{ IFF } \exists B \in P_f . \{s, t\} \subseteq B$$

IF:

- Hyp: $s \sim t$ Th: $\exists B \in P_f . \{s, t\} \subseteq B$
- By contradiction, s and t are in different blocks of the final partition
- At the beginning of the algorithm, P had a single block containing all the states
- s and t have been divided in a certain step of the algorithm, say the $(i + 1)$ -th
 - at step i , there was a partition $P_i = \{B_1, B_2, \dots, B_k\}$ such that $\exists j : 1 \leq j \leq k . \{s, t\} \subseteq B_j$
 - $P_{i+1} = \{B_1, B_2, \dots, B_j', B_j'', \dots, B_k\}$ with $s \in B_j'$ and $t \in B_j''$
- If s and t have been separated, $\exists \alpha . \text{tgt}_\alpha(s) \neq \text{tgt}_\alpha(t)$
 - $\exists s' : s \xrightarrow{\alpha} s' \wedge \forall t' (t \xrightarrow{\alpha} t' \vee \text{block}(s') \neq \text{block}(t'))$
 - $\exists t' : t \xrightarrow{\alpha} t' \wedge \forall s' (s \xrightarrow{\alpha} s' \vee \text{block}(t') \neq \text{block}(s'))$
- This contradicts $s \sim t$

ONLY IF:

- Hyp: $\exists B \in P_f . \{s,t\} \subseteq B$ Th: $s \sim t$
- Define $R = \{(x, y) \in B \times B : \exists B \in P_f\}$
- we prove that R is a bisimulation (notice that $\{s,t\}$ belongs to R by construction)
- It is symmetric, so it is enough that it is a simulation
- Let $(x,y) \in R$ and $x \xrightarrow{\alpha} x'$
- The block of x is no more splittable (the algorithm has terminated)
 - $\rightarrow \text{tgt}_\alpha(x) = \text{tgt}_\alpha(y)$
 - $\rightarrow \exists y' : y \xrightarrow{\alpha} y'$ and $\text{block}(x') = \text{block}(y')$
 - $\rightarrow x'$ and y' belong to the same block in P_f
 - \rightarrow By construction, $(x', y') \in R$

Complexity

- Let $(S, \Lambda, \rightarrow)$ be an LTS with $|S| = n$ and $|\rightarrow| = m$
- The final loop costs $O(n)$
- The loop *LOOP* is executed every time there is a partitioning
→ at most $O(n)$ times
- For the nested for-loops:
 - One loop for every block $B \in P$
 - One loop for every label $\alpha \in \Lambda$
 - The check $I \neq B$ can be performed in constant time, so as operation
 $P \leftarrow (P \setminus \{B\}) \cup I$
 - So the real cost is due to function *split*

```
P ← {S}
LOOP: forall B ∈ P do
    for all α ∈ Λ do
        I ← split(B, α, P)
        if I ≠ {B} then
            P ← (P \ {B}) ∪ I
            goto LOOP
forall B ∈ P do
    if {s,t} ⊆ B then return true
return false
```

- One loop for every state t belonging to the chosen block B
- The check $\text{tgt}_\alpha(s) = \text{tgt}_\alpha(t)$ performed for the chosen states s and t
 - We shall prove that this check can be performed by analyzing every α -trans. from states belonging to B only once
- So, the overall cost of the nested for-loops is

```

s ∈ B
B1 ← ∅  B2 ← ∅
forall t ∈ B do
    if tgtα(s) = tgtα(t) then B1 ← B1 ∪ {t}
    else B2 ← B2 ∪ {t}
if B2 = ∅ then return {B}
else return {B1, B2}

```

$$\sum_{B \in \mathcal{P}} \sum_{\alpha \in \Lambda} \sum_{t \in B} |\{t' : t \xrightarrow{\alpha} t'\}| = \sum_{B \in \mathcal{P}} \sum_{t \in B} \sum_{\alpha \in \Lambda} |\{t' : t \xrightarrow{\alpha} t'\}| = \sum_{t \in S} \sum_{\alpha \in \Lambda} |\{t' : t \xrightarrow{\alpha} t'\}| = |\rightarrow| = m$$

- Thus, the total cost is $O(nm)$

On efficiently handling blocks

- Symbolic representation of blocks
- Every block is a list of states and is associated to a binary string (+ lexicographic ordering)
- The starting block contains all states and has index 0
- When a block is split, the resulting sub-blocks are associated to the string of the original block with a new final bit, respectively 0 and 1

Example:

- let $S = \{1, 2, 3, 4, 5\}$ have index 0
- the first split produces ($B1 = \{1, 2, 4\}$, $B2 = \{3, 5\}$)
- Their associated binary strings will be 00 (for B1) and 01 (for B2)
- If then B2 is split into ($B3 = \{3\}$, $B4 = \{5\}$), then the associated strings will be 010 (for B3) and 011 (for B4)

- Then, we can consider the transition relation as a collection of triples (s, α, σ) , whenever $s \xrightarrow{\alpha} s'$, $s' \in B$ and σ is the binary index of B
- We fix an ordering on states and actions and lexicographically order the triples above
 - At the beginning this costs $O(m \log m)$
 - Every update of a block entails updating the triples that lead to that block (REMARK: this does NOT require a new ordering!)
- We obtain a matrix $(m \times 3)$ that looks like this:

Example:

$\begin{pmatrix} s_1 & \alpha_{1,1} & \sigma_{1,1} \\ s_1 & \alpha_{1,2} & \sigma_{1,2} \\ \dots & & \\ s_1 & \alpha_{1,k1} & \sigma_{1,k1} \\ s_2 & \alpha_{2,1} & \sigma_{2,1} \\ \dots & & \\ s_2 & \alpha_{2,k2} & \sigma_{2,k2} \\ \dots & & \\ s_n & \alpha_{n,1} & \sigma_{n,1} \\ \dots & & \\ s_n & \alpha_{n,kn} & \sigma_{n,kn} \end{pmatrix}$	$\begin{pmatrix} \dots & & \\ s_1 & \alpha_1 & 00 \\ s_1 & \alpha_2 & 01 \\ s_1 & \alpha_2 & 11 \\ \dots & & \end{pmatrix}$
	<p>If we split 01 (into 010 and 011), wherever s_1 has to go, the matrix will be</p> $\begin{pmatrix} \dots & & \\ s_1 & \alpha_1 & 00 \\ s_1 & \alpha_2 & 01b \\ s_1 & \alpha_2 & 11 \end{pmatrix}$

If in $\text{Split}(B, \alpha, P)$ we fix s as pivot state, then the check $\text{tgt}_\alpha(s) = \text{tgt}_\alpha(t)$ (for all t in B) is done as follows:

- Let k be the number of α -transitions from s
- If $k=0$, then
 - $B1 \leftarrow \{x \in B : x \text{ cannot perform } \alpha\}$
 - $B2 \leftarrow \{x \in B : x \text{ can perform } \alpha\}$
- Otherwise, let σ be the index of the first α -trans for s ; For every $t \in B$
 - if the index of the first α -trans for t is σ
 - then $B1 \leftarrow B1 \cup \{t\}$ else $B2 \leftarrow B2 \cup \{t\}$
- Let σ' be the index of the second α -trans for s ; For every $t \in B1$
 - if t has not a second α -trans or the index of its second α -trans is not σ'
 - then $B1 \leftarrow B1 \setminus \{t\}$ and $B2 \leftarrow B2 \cup \{t\}$
- ...
- Let σ'' be the index of the k th α -trans for s ; For every $t \in B1$
 - if t has not a k th α -trans or the index of its k th α -trans is not σ''
 - then $B1 \leftarrow B1 \setminus \{t\}$ and $B2 \leftarrow B2 \cup \{t\}$
- For every $t \in B1$
 - if t has other α -trans
 - then $B1 \leftarrow B1 \setminus \{t\}$ and $B2 \leftarrow B2 \cup \{t\}$