Logical Characterisations of Probabilistic Bisimilarity

Yuxin Deng

East China Normal University

(Based on joint work with Hengyang Wu and Yuan Feng)

IFIP Working Group 2.2 meeting, Bordeaux, September 18, 2017

Preliminaries

Labelled transition systems

Def. A *labelled transition system* (LTS) is a triple $\langle S, Act, \rightarrow \rangle$, where

1. S is a set of states
2. Act is a set of actions
3. $\rightarrow \;\subseteq\; S \times Act \times S$ is the transition relation

Write $s \xrightarrow{\alpha} s'$ for $(s, \alpha, s') \in \rightarrow$.

Bisimulation

$$\begin{array}{ccc} s & \xrightarrow{a} & s' \\ \mathcal{R} & & \mathcal{R} \\ t & \xrightarrow{a} & t' \end{array}$$

s and t are bisimilar if there exists a bisimulation $\mathcal{R}$ with $s \mathrel{\mathcal{R}} t$.

Probabilistic labelled transition systems

Def. A *probabilistic labelled transition system* (pLTS) is a triple $\langle S, Act, \rightarrow \rangle$, where

1. S is a set of states
2. Act is a set of actions
3. $\rightarrow \;\subseteq\; S \times Act \times \mathcal{D}(S)$.

We usually write $s \xrightarrow{\alpha} \Delta$ in place of $(s, \alpha, \Delta) \in \rightarrow$.

Example

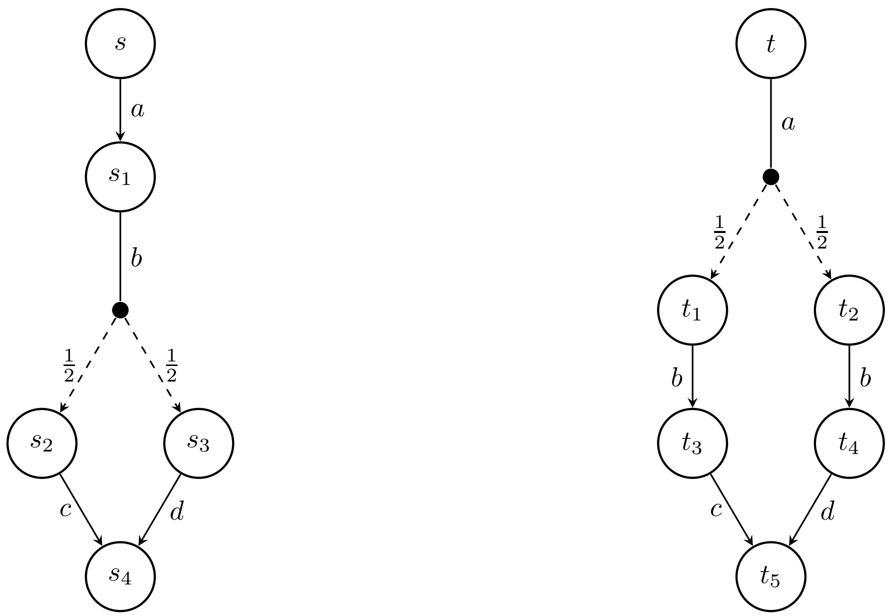

Probabilistic Bisimulation

$$\begin{array}{ccc} s & \xrightarrow{a} & \Delta \\ \mathcal{R} & & \mathcal{R}^{\dagger} \\ t & \xrightarrow{a} & \Theta \end{array}$$

Write $\sim$ for probabilistic bisimilarity.

Lifting relations

Def. Let S, T be two countable sets and $\mathcal{R} \subseteq S \times T$ be a binary relation. The lifted relation $\mathcal{R}^\dagger \subseteq \mathcal{D}(S) \times \mathcal{D}(T)$ is the smallest relation satisfying

1. $s \; \mathcal{R} \; t$ implies $\overline{s} \mathcal{R}^\dagger \overline{t}$
2. $\Delta_i \mathcal{R}^\dagger \Theta_i$ for all $i \in I$ implies $(\sum_{i \in I} p_i \cdot \Delta_i) \mathcal{R}^\dagger (\sum_{i \in I} p_i \cdot \Theta_i)$

There are alternative formulations; related to the Kantorovich metric and the network flow problem. See e.g. `http://www.springer.com/978-3-662-45197-7`

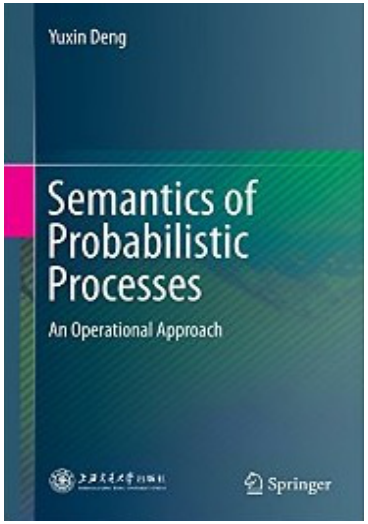

The first modal characterisation

The logic $\mathcal{L}_1$

The language $\mathcal{L}_1$ of formulas:

$$\varphi ::= \quad \top \mid \varphi_1 \wedge \varphi_2 \mid \langle a \rangle_p \varphi.$$

where p is rational number in $[0, 1]$.

Semantics

- $s \models \top$ always;
- $s \models \varphi_1 \wedge \varphi_2$, if $s \models \varphi_1$ and $s \models \varphi_2$;
- $s \models \langle a \rangle_p \varphi$ iff $s \xrightarrow{a} \Delta$ and $\Delta([\![\varphi]\!]) \geq p$, where $[\![\varphi]\!] = \{s \in S \mid s \models \varphi\}$.

Logical equivalence: $s =_1 t$ if $s \models \varphi \Leftrightarrow t \models \varphi$ for all $\varphi \in \mathcal{L}_1$.

Modal characterisation

Modal characterisation ($s \sim t$ iff $s =_1 t$) for the continuous case given by [Desharnais et al. Inf. Comput. 2003], using the machinery of analytic spaces.

298 *Prakash Panangaden*

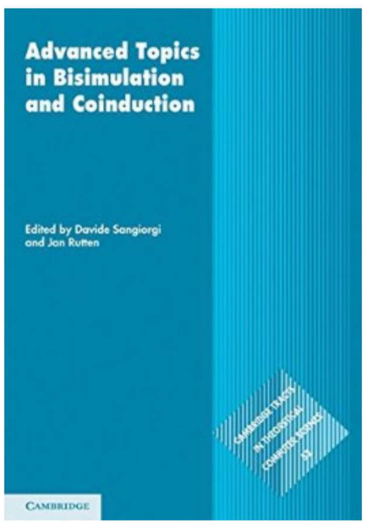

There are many variations that one can imagine. Perhaps the simplest is to have negation and dispense with Δ and disjunction. All the variations considered by Larsen and Skou have some negative construct. The striking fact – first discovered in the context of continuous state spaces [DEP98, DEP02] – is that one can get a logical characterisation result with purely positive formulas. The discrete case is covered by these results. Surprisingly no elementary proof for the discrete case – i.e. one that avoids the measure theory machinery – is known.

The π-λ theorem

Let $\mathcal{P}$ be a family of subsets of a set X. $\mathcal{P}$ is a π-class if it is closed under finite intersection; $\mathcal{P}$ is a λ-class if it is closed under complementations and countable disjoint unions.

Thm. If $\mathcal{P}$ is a π-class, then $\sigma(\mathcal{P})$ is the smallest λ-class containing $\mathcal{P}$, where $\sigma(\mathcal{P})$ is a σ-algebra containing $\mathcal{P}$.

An application of the π-λ theorem

Prop. Let $\mathcal{A}_0 = \{[\![\varphi]\!] \mid \varphi \in \mathcal{L}\}$. For any $\Delta, \Theta \in \mathcal{D}(S)$, if $\Delta(A) = \Theta(A)$ for any $A \in \mathcal{A}_0$, then $\Delta(B) = \Theta(B)$ for any $B \in \sigma(\mathcal{A}_0)$.

Soundness and completeness of the logic

Lem. Given the logic $\mathcal{L}$, and let $(S, A, \rightarrow)$ be a reactive pLTS with countably many states. Then for any two states $s, t \in S$, $s \sim t$ iff $s =_1 t$.

Proof. Use the π-λ theorem. See [Deng and Wu. ICFEM 2014].

The second modal characterisation

The logic $\mathcal{L}_2$

The language $\mathcal{L}_2$ of formulas:

$$\varphi ::= \quad \top \mid \varphi_1 \wedge \varphi_2 \mid \langle a \rangle \varphi.$$

Modal characterisation for the continuous case given by [van Breugel et al. TCS 2005], using the machinery of probabilistic powerdomains and Banach algebra.

We will see the discrete case can be much simplified.

Semantics

$$
\begin{array}{rcl}
Pr(s, \top) & = & 1 \\
Pr(s, \langle a \rangle \varphi) & = & \begin{cases} \sum_{t \in \lceil \Delta \rceil} \Delta(t) \cdot Pr(t, \varphi) & \text{if } s \xrightarrow{a} \Delta \\ 0 & \text{otherwise.} \end{cases} \\
Pr(s, \varphi_1 \wedge \varphi_2) & = & Pr(s, \varphi_1) \cdot Pr(s, \varphi_2)
\end{array}
$$

Logical equivalence: $s =_2 t$ if $Pr(s, \varphi) = Pr(t, \varphi)$ for all $\varphi \in \mathcal{L}_2$.

Soundness

Thm. If $s \sim t$ then $s =_2 t$.

Proof. Easy by structural induction.

Completeness

Thm. For finite-state reactive pLTSs, if $s =_2 t$ then $s \sim t$.

Proof.

- Observe that $=_2$ is an equivalence relation.
- Let $C_1, C_2, ..., C_n$ be all the equivalence classes.
- Write $Pr(C_i, \varphi)$ for $Pr(s_{ij,\varphi})$, where $s_{ij} \in C_i$ and $\varphi \in \mathcal{L}_2$.
- For any $i \neq j$, let φ_{ij} be a distinguishing formula with $Pr(C_i, \varphi_{ij}) \neq Pr(C_j, \varphi_{ij})$.

Key lemma

Lem. For any $I \subseteq \{1, \cdots, n\}$ with $I \neq \emptyset$, there exist a nonempty $I' \subseteq I$ and an enhanced formula φ such that

(i) for any $i \in I$, $i \in I'$ iff $Pr(C_i, \varphi) > 0$;

(ii) for any $i \neq j \in I'$, $Pr(C_i, \varphi) \neq Pr(C_j, \varphi)$.

Algorithm for computing enhanced formulas

---

input : A nonempty subset I of $\{1, \cdots, n\}$ with the distinguishing formula φ_{ij} for all $i \neq j$.

output: A nonempty $I' \subseteq I$ and an enhanced formula φ satisfying (i) and (ii) in the key lemma.

begin
- $\mathcal{I}_{pass} \leftarrow \emptyset$; $\mathcal{I}_{rem} \leftarrow \{(i,j) \in I \times I : i < j\}$; $I' \leftarrow I$; $\varphi \leftarrow \top$;
- **while** $\mathcal{I}_{rem} \neq \emptyset$ **do**
  - Choose arbitrarily $(i,j) \in \mathcal{I}_{rem}$;
  - $I' \leftarrow \{k \in I' : Pr(C_k, \varphi_{ij}) > 0\}$;
  - $\mathcal{I}_{dis} \leftarrow \{(k,l) \in \mathcal{I}_{rem} \cap I' \times I' : Pr(C_k, \varphi_{ij}) \neq Pr(C_l, \varphi_{ij})\}$;
  - $\mathcal{I}_{rem} \leftarrow (\mathcal{I}_{rem} \cap I' \times I') \backslash \mathcal{I}_{dis}$; $\mathcal{I}_{pass} \leftarrow (\mathcal{I}_{pass} \cap I' \times I') \cup \mathcal{I}_{dis}$; $\varphi \leftarrow \varphi \wedge \varphi_{ij}$;
  - $\mathcal{I}_{tem} \leftarrow \emptyset$; $\mathcal{I} \leftarrow \mathcal{I}_{pass}$;
  - **while** $\mathcal{I} \neq \emptyset$ **do**
    - $\mathcal{I} \leftarrow \{(k,l) \in \mathcal{I}_{pass} \backslash \mathcal{I}_{tem} : Pr(C_k, \varphi) = Pr(C_l, \varphi)\}$;
    - **if** $\mathcal{I} \neq \emptyset$ **then**
      - $\varphi \leftarrow \varphi \wedge \varphi_{ij}$; $\mathcal{I}_{tem} \leftarrow \mathcal{I}_{tem} \cup \mathcal{I}$;
    - **end**
  - **end**
- **end**
- **return** I', φ;

end

Correctness of the algorithm

The algorithm has recently been formalized in Coq. Correctness proof relies on four invariants of the outer loop:

(a) $I' \neq \emptyset$;

(b) for any $i \in I$, $\quad i \in I'$ iff $Pr(C_i, t) > 0$;

(c) $\mathcal{I}_{pass} \cup \mathcal{I}_{rem} = \{(i,j) \in I' \times I' : i < j\}$;

(d) for any $(i,j) \in \mathcal{I}_{pass}$, $Pr(C_i, t) \neq Pr(C_j, t)$.

Non-trivial proofs at all, with about 1500 lines of Coq code used.

Completeness proof

- Suppose $s =_2 t$. A transition $s \xrightarrow{a} \Delta$ has to be matched by $t \xrightarrow{a} \Theta$. It remains to show $\Delta (=_2)^\dagger \Theta$.
- It suffices to show $\Delta(C_i) = \Theta(C_i)$ for all equivalence classes C_i with $i \in I$.
- By induction on $|I|$. The case $|I| = 1$ trivial.
- Let φ be any formula.

$$0 = Pr(s, \langle a \rangle \varphi) - Pr(t, \langle a \rangle \varphi) = \sum_{i \in I} Pr(C_i, \varphi) \cdot (\Delta(C_i) - \Theta(C_i))$$

- The key lemma gives some $I' \subseteq I$ and enhanced formula φ_0. Let $a_i = Pr(C_i, \varphi_0)$ and $x_i = \Delta(C_i) - \Theta(C_i)$.
- Then $a_1 x_1 + a_2 x_2 + \cdots + a_n x_n = 0$, where $I' = \{1, ..., n\}$.

- Any formula $\wedge^m \varphi_0$ gives the equation $a_1^m x_1 + a_2^m x_2 + \cdots + a_n^m x_n = 0$.
-

$$\begin{aligned} a_1 x_1 + a_2 x_2 + \cdots + a_n x_n &= 0 \\ a_1^2 x_1 + a_2^2 x_2 + \cdots + a_n^2 x_n &= 0 \\ &\vdots \\ a_1^n x_1 + a_2^n x_2 + \cdots + a_n^n x_n &= 0 \end{aligned}$$

- Modify the coefficient matrix to get

$$\begin{bmatrix} 1 & 1 & 1 & \cdots & 1 \\ a_1 & a_2 & a_3 & \cdots & a_n \\ a_1^2 & a_2^2 & a_3^2 & \cdots & a_n^2 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ a_1^{n-1} & a_2^{n-1} & a_3^{n-1} & \cdots & a_n^{n-1} \end{bmatrix}$$

— the transpose of a Vandermonde matrix.

- $x_i = 0$, i.e., $\Delta(C_i) = \Theta(C_i)$ for all $i \in I'$.
- $\sum_{i \in I \setminus I'} Pr(C_i, \varphi) \cdot (\Delta(C_i) - \Theta(C_i)) = 0$
- $|I \setminus I'| < |I|$ and by induction we get $\Delta(C_i) = \Theta(C_i)$ for all $i \in I \setminus I'$.
- $\Delta(=_2)^{\dagger}\Theta$ as required.

Summary

Two logical characterisation of probabilistic bisimilarity for countable and finite-state reactive processes, respectively, with much simpler proofs than those of Desharnais et al. and van Breugel et al.

Thank you!