Disjunctive Probabilistic Modal Logic is Enough for Bisimilarity on Reactive Probabilistic Systems

Marco Bernardo

University of Urbino – Italy

joint work with: Marino Miculan

Process Model

- A reactive probabilistic labeled transition system (RPLTS) is a triple $(S, A, \longrightarrow)$ where:
  - S is a countable set of *states*.
  - A is a countable set of *actions*.
  - $\longrightarrow \subseteq S \times A \times \mathit{Distr}(S)$ is a *transition relation* such that, whenever $s \xrightarrow{a} \Delta_1$ and $s \xrightarrow{a} \Delta_2$, then $\Delta_1 = \Delta_2$.
- No internal nondeterminism.
- Rabin probabilistic automata.
- $\mathit{Distr}(S)$ is the set of discrete probability distributions over S having *finite support*: $\mathit{supp}(\Delta) \triangleq \{s \in S \mid \Delta(s) > 0\}$ is finite.
- Probabilistic counterpart of image finiteness.

Process Semantics

- Larsen & Skou defined probabilistic bisimilarity $\sim_{\mathrm{PB}}$ over RPLTS.
- An equivalence relation $\mathcal{B}$ over S is a probabilistic bisimulation iff, whenever $(s_1, s_2) \in \mathcal{B}$, then for all actions $a \in A$:
  - If $s_1 \xrightarrow{a} \Delta_1$ then there exists $s_2 \xrightarrow{a} \Delta_2$ such that $\Delta_2(C) = \Delta_1(C)$ for all equivalence classes $C \in S/\mathcal{B}$.
  - If $s_2 \xrightarrow{a} \Delta_2$ then there exists $s_1 \xrightarrow{a} \Delta_1$ such that $\Delta_1(C) = \Delta_2(C)$ for all equivalence classes $C \in S/\mathcal{B}$.
- States $s_1, s_2 \in S$ are probabilistically bisimilar, written $s_1 \sim_{\mathrm{PB}} s_2$, iff there exists a probabilistic bisimulation including (s_1, s_2).
- Probabilistic extension of Milner strong bisimilarity.
- Ordinary lumpability for (discrete-time) Markov chains.

Process Logic

- Larsen & Skou defined probabilistic modal logic PML over RPLTS.
- Probabilistic extension of Hennessy & Milner logic in which $\langle a \rangle$ is decorated with a probabilistic lower bound $p \in \mathbb{R}_{[0,1]}$.
- Variants of PML:

$$
\begin{array}{ll}
\text{PML}_{\neg\wedge} & \phi ::= \top \mid \neg\phi \mid \phi \wedge \phi \mid \langle a \rangle_p \phi \\
\text{PML}_{\neg\vee} & \phi ::= \top \mid \neg\phi \mid \phi \vee \phi \mid \langle a \rangle_p \phi \\
\text{PML}_{\wedge} & \phi ::= \top \mid \phi \wedge \phi \mid \langle a \rangle_p \phi \\
\text{PML}_{\vee} & \phi ::= \top \mid \phi \vee \phi \mid \langle a \rangle_p \phi
\end{array}
$$

where:

$s \models \langle a \rangle_p \phi$ iff there exists $s \xrightarrow{a} \Delta$ s.t. $\Delta(\{s' \in S \mid s' \models \phi\}) \geq p$

Logical Characterization of Behavioral Equivalences

- The equivalence induced by a logic (satisfying the same formulas) coincides with a given behavioral equivalence.
- Distinguishing formulas useful to explain inequivalence.
- Larsen & Skou showed that $\mathrm{PML}_{\neg\wedge}$ characterizes $\sim_{\mathrm{PB}}$.
- Minimal deviation assumption (stronger than finite support).
- Desharnais, Edalat & Panangaden showed that also $\mathrm{PML}_{\wedge}$ characterizes $\sim_{\mathrm{PB}}$, i.e., *negation* is not necessary over RPLTS.
- No assumptions, continuous state spaces, measure theory.
- Jacobs and Sokolova proved it again coalgebraically, in the setting of dual adjunctions between spaces and logics.

Summary of Results

- $\mathrm{PML}_{\vee}$ characterizes $\sim_{\mathrm{PB}}$, i.e., disjunction is enough for reactive probabilistic processes.
- Proof based on a coalgebraic characterization of $\sim_{\mathrm{PB}}$ for finite reactive probabilistic trees.
- A simpler proof that $\mathrm{PML}_{\neg\wedge}$ characterizes $\sim_{\mathrm{PB}}$, which does not require the minimal deviation assumption.
- A simpler proof that $\mathrm{PML}_{\wedge}$ characterizes $\sim_{\mathrm{PB}}$, which works directly with discrete state spaces.
- Effective proofs showing how to build distinguishing formulas.

Reactive Probabilistic Trees

- Each RPLTS can be given a semantics in a canonical form, which we call reactive probabilistic trees.
- Categorical construction based on Rutten, de Vink, Worrell.
- Probabilistic counterpart of Winskel synchronization trees.
- *Fully abstract* semantics: $s_1 \sim_{\mathrm{PB}} s_2$ iff they are mapped to the *same* RPT t.
- *Compact* semantics: $t_1 = t_2$ iff all of their *finite* height approximations are pairwise equal.
- Summing up: $s_1 \sim_{\mathrm{PB}} s_2$ iff all the *finite* height approximations of the RPT for s_1 and the RPT for s_2 are pairwise equal.

Working with Finite RPT

- Finding a logical characterization of $\sim_{\mathrm{PB}}$ over RPLTS reduces to finding a logical characterization of $=$ over *finite* RPT.
- Build a distinguishing formula *by induction* on their *finite* height when $t_1 \neq t_2$, but be careful!
- An additional constraint has to be met.
- If a variant of PML characterizes $=$ over *finite* RPT and for any two nodes t_1 and t_2 of a *finite* RPT such that $t_1 \neq t_2$ there exists a formula ϕ distinguishing t_1 from t_2 such that:

$$depth(\phi) \leq \max(height(t_1), height(t_2))$$

then the variant of PML characterizes $\sim_{\mathrm{PB}}$ over RPLTS.

Depth of Formulas and Height of Trees

- If $depth(\phi)$ were greater, then:
  - ϕ may not distinguish higher finite approximations of s_1 and s_2;
  - no shorter formula derivable from ϕ may still distinguish t_1 and t_2.
- $s_1 \xrightarrow{a} s_1' \xrightarrow{c} s_1'' \xrightarrow{e} s_1'''$ and $s_2 \xrightarrow{b} s_2' \xrightarrow{d} s_2'' \xrightarrow{f} s_2'''$ differ at $height = 1$, so we can focus on $t_1 \xrightarrow{a} t_1'$ and $t_2 \xrightarrow{b} t_2'$.
- $\phi = \langle a \rangle_1 \neg \langle c \rangle_1$, of $depth = 2$, distinguishes $t_1 \xrightarrow{a} t_1'$ and $t_2 \xrightarrow{b} t_2'$, but does *not* distinguish $t_1 \xrightarrow{a} t_1' \xrightarrow{c} t_1''$ and $t_2 \xrightarrow{b} t_2' \xrightarrow{d} t_2''$.
- $\phi = \langle a \rangle_1 \vee \langle b \rangle_1 \langle c \rangle_1$, of $depth = 2$, tells apart $t_1 \xrightarrow{a} t_1'$ and $t_2 \xrightarrow{b} t_2'$, but the derived shorter formula $\langle a \rangle_1 \vee \langle b \rangle_1$ of $depth = 1$ does *not*.

A New Proof that $\mathrm{PML}_{\neg\wedge}$ Characterizes $\sim_{\mathrm{PB}}$

- Given $t_1 \neq t_2$, if one has an action a not possessed by the other, then $\langle a \rangle_1$ tells them apart.
- If they have the same actions, then there must exist an action a such that $t_1 \xrightarrow{a} \Delta_{1,a}$ and $t_2 \xrightarrow{a} \Delta_{2,a}$ with $\Delta_{1,a} \neq \Delta_{2,a}$.
- Consider $t' \in \mathit{supp}(\Delta_{1,a})$ such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$.
- Let $\mathit{supp}(\Delta_{2,a}) \backslash \{t'\} = \{t'_{2,1}, t'_{2,2}, \ldots, t'_{2,k}\}$, which cannot be empty.
- For each $j = 1, 2, \ldots, k$, by the induction hypothesis there exists $\phi'_{2,j} \in \mathrm{PML}_{\neg\wedge}$ meeting $\mathit{depth}(\phi'_{2,j}) \leq \max(\mathit{height}(t'), \mathit{height}(t'_{2,j}))$ such that $t' \models \phi'_{2,j} \not\models t'_{2,j}$.
- We can impose direction of $\phi'_{2,j}$-satisfaction thanks to negation!
- Therefore:

$$t_1 \models \langle a \rangle_{\Delta_{1,a}(t')} \bigwedge_{1 \leq j \leq k} \phi'_{2,j} \not\models t_2$$

Proving that $\mathrm{PML}_{\neg\vee}$ Characterizes $\sim_{\mathrm{PB}}$

- $\mathrm{PML}_{\neg\vee}$ is obviously equivalent to $\mathrm{PML}_{\neg\wedge}$ due to De Morgan laws.
- Useful intermediate step to achieve our result for $\mathrm{PML}_{\vee}$.
- Adaptation of the proof for $\mathrm{PML}_{\neg\wedge}$.
- Recall that $t' \in \mathit{supp}(\Delta_{1,a})$ is such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$ and that $\mathit{supp}(\Delta_{2,a}) \backslash \{t'\} = \{t'_{2,1}, t'_{2,2}, \ldots, t'_{2,k}\}$ is not empty.
- For each $j = 1, 2, \ldots, k$, by the induction hypothesis there exists $\phi'_{2,j} \in \mathrm{PML}_{\neg\vee}$ meeting $\mathit{depth}(\phi'_{2,j}) \leq \max(\mathit{height}(t'), \mathit{height}(t'_{2,j}))$ such that $t' \not\models \phi'_{2,j} =\!\!| \, t'_{2,j}$.
- We can impose direction of $\phi'_{2,j}$-satisfaction thanks to negation!
- Therefore:

$$t_1 \not\models \langle a \rangle_{1-\Delta_{2,a}(t')} \bigvee_{1 \leq j \leq k} \phi'_{2,j} =\!\!| \, t_2$$

Proving that $\mathrm{PML}_{\vee}$ Characterizes $\sim_{\mathrm{PB}}$

- Negation is no longer available!
- However, the proof for $\mathrm{PML}_{\neg\vee}$ is still useful.
- The objective is to get to:

$$t_1 \not\models \langle a \rangle_{1-(\Delta_{2,a}(t')+p)} \bigvee_{j \in J} \phi'_{2,j} \not\dashv t_2$$

  for:

  - a derivative t' of t_1 s.t. $\Delta_{1,a}(t') > \Delta_{2,a}(t')$ and $t' \not\models \phi'_{2,j} \not\dashv t'_{2,j}$;
  - a suitable probabilistic value p such that $\Delta_{2,a}(t') + p < 1$;
  - an index set J identifying certain derivatives of t_2 other than t'.

- The choice of t' is crucial!
- Intuition: t' has to satisfy as few $\mathrm{PML}_{\vee}$ formulas as possible.

From $\wedge$ to $\vee$ within Distinguishing Formulas

- A disjunctive distinguishing formula can often be obtained from a conjunctive distinguishing formula by *increasing* in the latter some of its probabilistic lower bounds.
- Example:

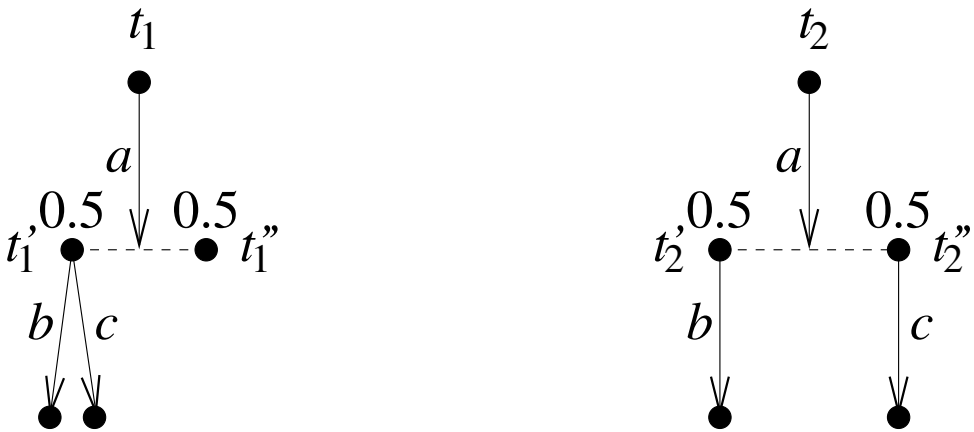

$$t_1 \models \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1) \not\models t_2$$

$$t_1 \not\models \langle a \rangle_{1.0} (\langle b \rangle_1 \vee \langle c \rangle_1) \models t_2$$

Direction of Distinguishing Formula Satisfaction

- It is not always the case that the direction of distinguishing formula satisfaction is inverted when moving from $\wedge$ to $\vee$.
- Example:

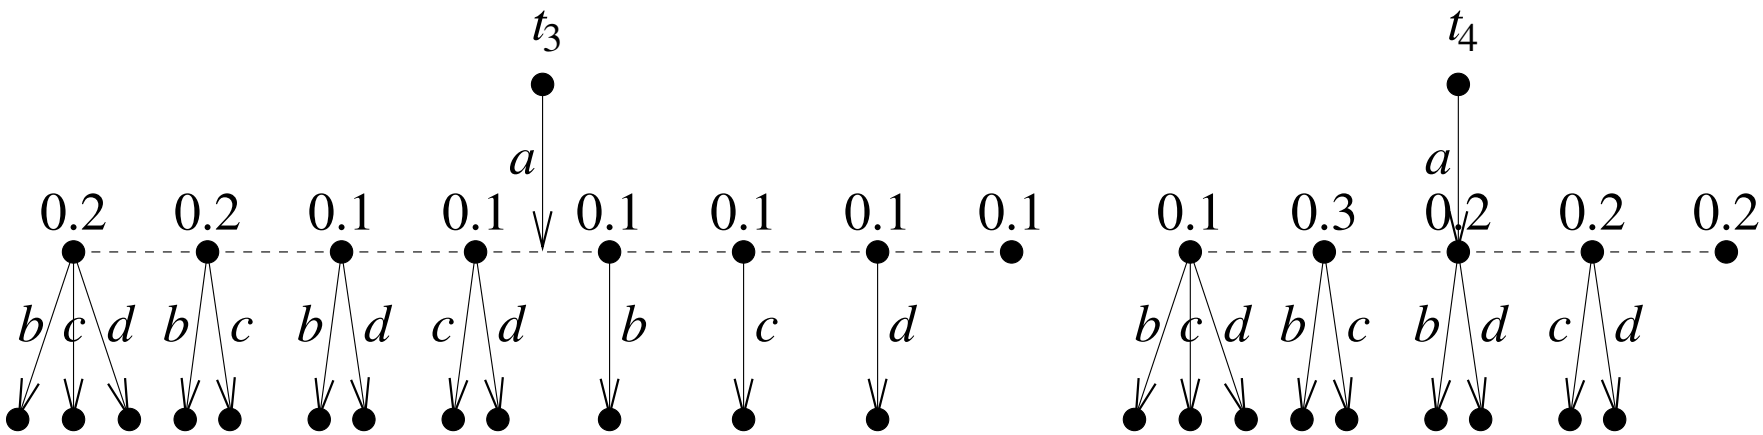

$$t_3 \models \langle a \rangle_{0.2} (\langle b \rangle_1 \wedge \langle c \rangle_1 \wedge \langle d \rangle_1) \not\models t_4$$

$$t_3 \models \langle a \rangle_{0.9} (\langle b \rangle_1 \vee \langle c \rangle_1 \vee \langle d \rangle_1) \not\models t_4$$

No Increase when $\wedge$ and $\vee$ Are Not Necessary

- Increasing some of the probabilistic lower bounds in a conjunctive distinguishing formula does not always yield a disjunctive one, especially when $\wedge$ and $\vee$ are not necessary for discriminating.
- Example:

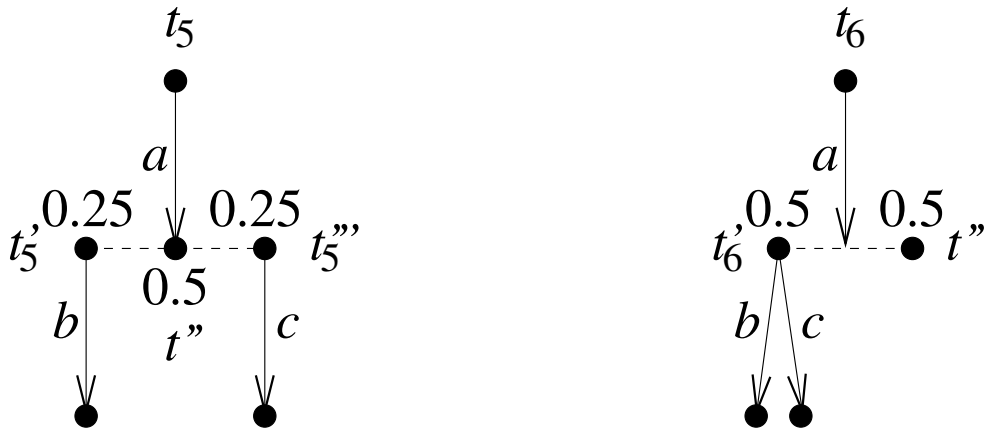

$$t_5 \not\models \langle a \rangle_{0.5} (\langle b \rangle_1 \wedge \langle c \rangle_1) \not\dashv t_6$$

$$t_5 \not\models \langle a \rangle_{0.5} \langle b \rangle_1 \not\dashv t_6$$

Set of $\text{PML}_\vee$ Formulas Satisfied by a Finite RPT Node

- For choosing t' in $\langle a \rangle_{1-(\Delta_{2,a}(t')+p)} \bigvee_{j \in J} \phi'_{2,j}$, we need to know the set of $\text{PML}_\vee$ formulas that are satisfied by each node t.
- Consider only the $\text{PML}_\vee$ formulas satisfied by t featuring:
  - probabilistic lower bounds of diamonds that are *maximal* with respect to the satisfiability of a formula of that form by t (so to keep the set finite);
  - diamonds that arise only from *existing* transitions departing from t (to avoid useless diamonds in disjunctions and keep the set finite);
  - disjunctions that:
    - stem only from *single* transitions of *different* nodes in the support of a distribution reached by t;
    - are preceded by a diamond decorated with the *sum* of the probs assigned to those nodes by the distribution reached by t.

Formalization of the $\Phi_\vee$-Set

- If $height(t) = 0$, then $\Phi_\vee(t) = \emptyset$.
- If $height(t) \geq 1$ for t having transitions of the form $t \xrightarrow{a_i} \Delta_i$ with $supp(\Delta_i) = \{t'_{i,j} \mid j \in J_i\}$ and $i \in I \neq \emptyset$, then:

$$\begin{aligned}\Phi_\vee(t) \;=\; & \{\langle a_i \rangle_1 \mid i \in I\} \\ & \cup \bigcup_{i \in I} hplb\Big(\bigcup_{\emptyset \neq J' \subseteq J_i} \Big\{ \Big\langle a_i \Big\rangle_{\sum_{j \in J'} \Delta_i(t'_{i,j})} \dot{\bigvee_{j \in J'}} \phi'_{i,j,k} \;\Big|\; \\ & \qquad t'_{i,j} \in supp(\Delta_i), \phi'_{i,j,k} \in \Phi_\vee(t'_{i,j}) \Big\} \Big)\end{aligned}$$

where:

- $\dot{\vee}$ is a variant of $\vee$ in which identical operands are not admitted (i.e., idempotence is forced);
- $hplb$ keeps only the formula with the highest probabilistic lower bound decorating the initial a_i-diamond among the formulas differring only for that bound.

Selection of the Parameters

- Among the nodes such that $\Delta_{1,a}(t') > \Delta_{2,a}(t')$, avoid those for which there exists another node whose formulas *without* $\vee$ have the same form but *higher* probabilistic lower bounds.
- Then a good criterion for choosing t' in $\langle a \rangle_{1-(\Delta_{2,a}(t')+p)} \bigvee_{j \in J} \phi'_{2,j}$ is the *minimality* of its $\Phi_\vee$-set, so to obtain $t' \not\models \phi'_{2,j} \dashv t'_{2,j}$.
- The set J only contains the derivatives of t_2 different from t' to which $\Delta_{1,a}$ and $\Delta_{2,a}$ assign two *different* probabilities.
- The excluded derivatives do not matter for discriminating: some of them satisfy $\bigvee_{j \in J} \phi'_{2,j}$, the others do not.
- The value p is the probability that t_2 (same as for t_1) reaches the *excluded* derivatives that do *not* satisfy $\bigvee_{j \in J} \phi'_{2,j}$.

A New Proof that $\mathrm{PML}_{\wedge}$ Characterizes $\sim_{\mathrm{PB}}$

- Adaptation of the proof for $\mathrm{PML}_{\vee}$ based on the one for $\mathrm{PML}_{\neg\wedge}$.
- The objective is to get to:

$$t_1 \models \langle a \rangle_{\Delta_{1,a}(t')+p} \bigwedge_{j \in J} \phi'_{2,j} \not\models t_2$$

  for:

  - a derivative t' of t_1 s.t. $\Delta_{1,a}(t') > \Delta_{2,a}(t')$ and $t' \models \phi'_{2,j} \not\models t'_{2,j}$;
  - a suitable probabilistic value p such that $\Delta_{1,a}(t') + p \leq 1$;
  - an index set J identifying certain derivatives of t_2 other than t'.

- Construction of the $\Phi_{\wedge}$-set.
- Choose t' based on the *maximality* of its $\Phi_{\wedge}$-set.
- The value p is the probability that t_2 (same as for t_1) reaches the *excluded* derivatives that *satisfy* $\bigwedge_{j \in J} \phi'_{2,j}$.

Summary of Results

- $\mathrm{PML}_{\vee}$ characterizes $\sim_{\mathrm{PB}}$, i.e., disjunction is enough for reactive probabilistic processes.
- Proof based on a coalgebraic characterization of $\sim_{\mathrm{PB}}$ for finite reactive probabilistic trees.
- A simpler proof that $\mathrm{PML}_{\neg\wedge}$ characterizes $\sim_{\mathrm{PB}}$, which does not require the minimal deviation assumption.
- A simpler proof that $\mathrm{PML}_{\wedge}$ characterizes $\sim_{\mathrm{PB}}$, which works directly with discrete state spaces.
- Effective proofs showing how to build distinguishing formulas.