Part 3: Labelled Markov Processes

Prakash Panangaden
McGill University

What are LMPs?

- Labelled Markov processes are probabilistic versions of labelled transition systems.
- A transition is triggered by an action (label) and the final state is governed by a probability distribution that depends on the action and the initial state.
- There is no other indeterminacy.
- No probabilities are associated with the choice of label, i.e. no attempt to model the environment.

- We observe the actions, not the internal states. Unlike most of the treatments seen in probability books. This does not make a serious difference.

- **In general, the state space may be a continuum.**

Motivation

- Model and reason about systems with continuous state space or continuous time evolution or both.
- Hybrid systems
- Telecommunication systems with spatial variation, e.g. mobile phones
- performance modelling
- probabilistic languages with loops or recursion

An example

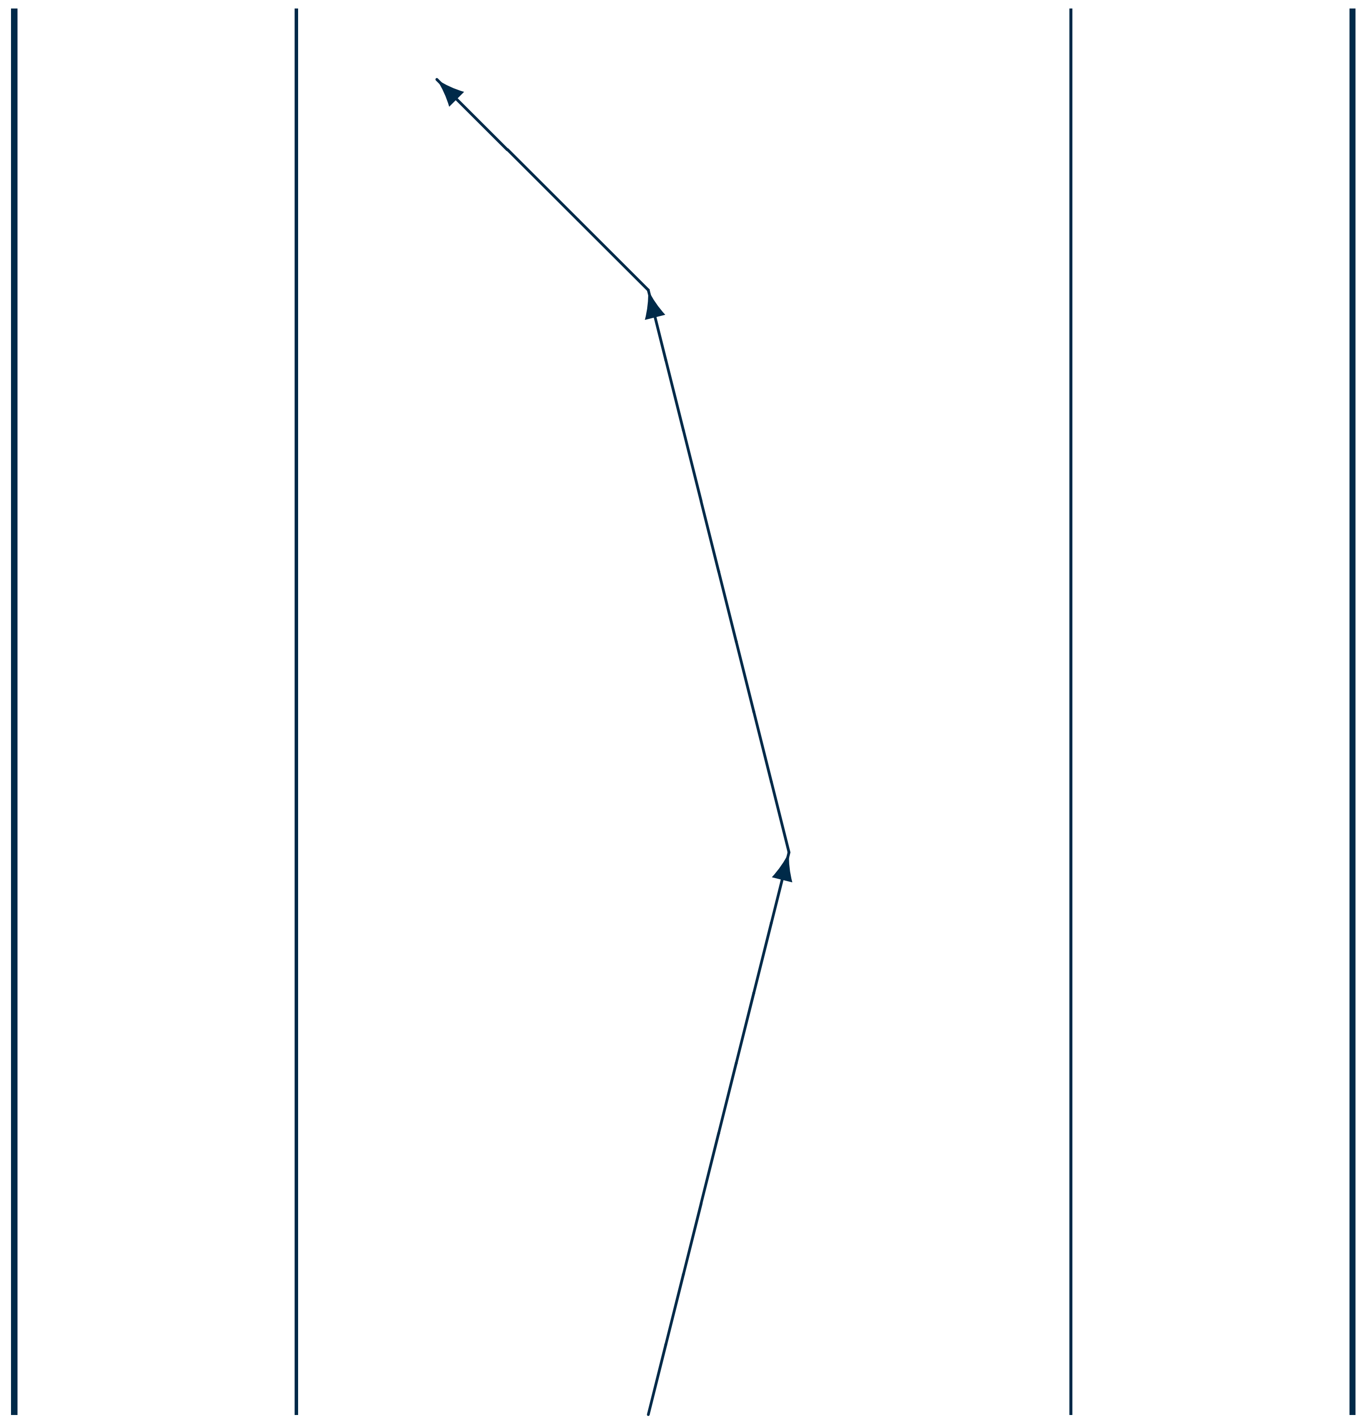

a - turn left

b - turn right

c - straight

a - turn left, b - turn right, c - keep on course

The actions move the craft sideways with some probability distributions on how far it moves. The craft may “drift” even with c. The action a (b) must be disabled when the craft is too near the left (right) boundary.

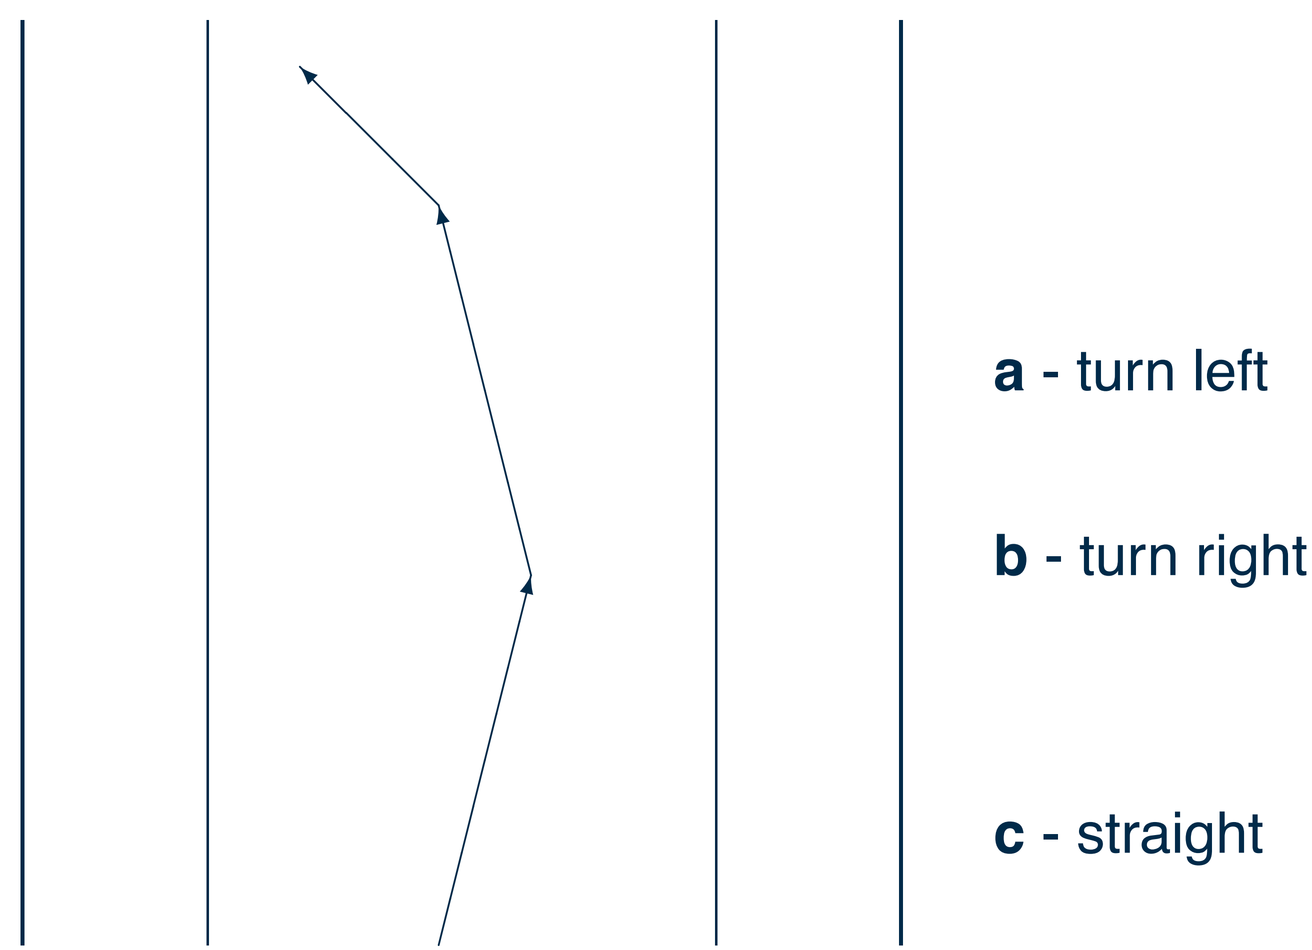

Schematic of the example

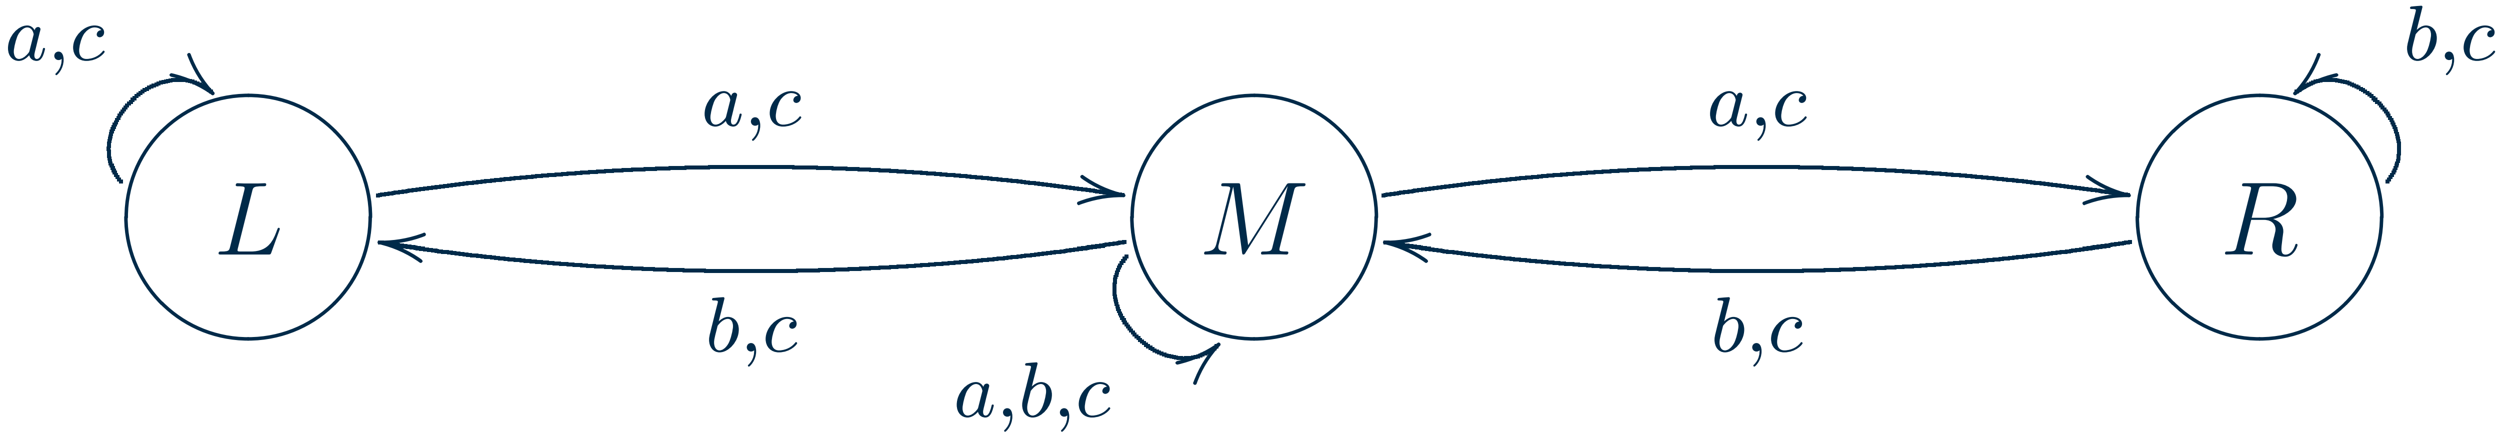

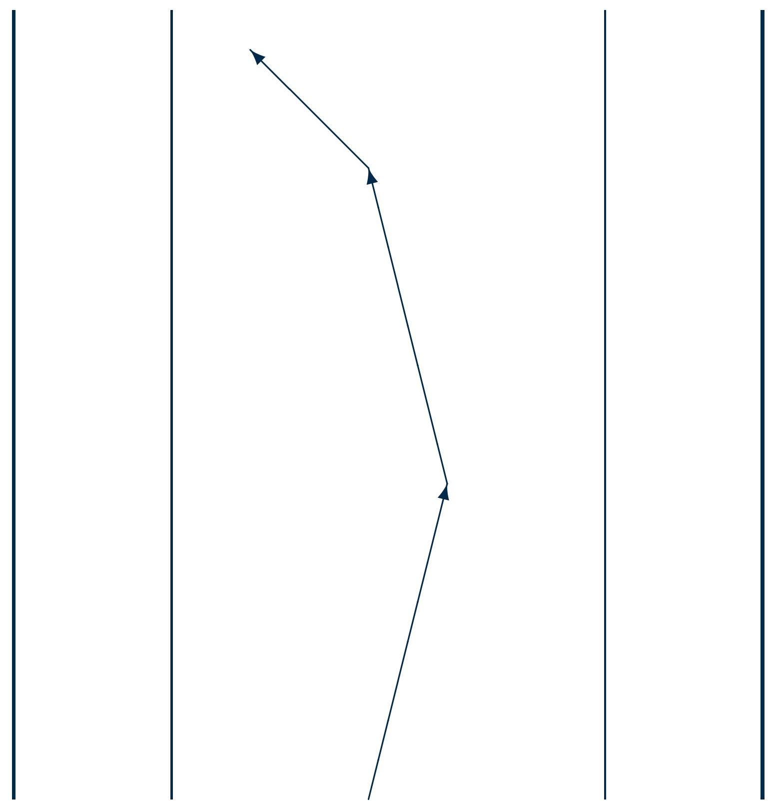

a - turn left

b - turn right

c - straight

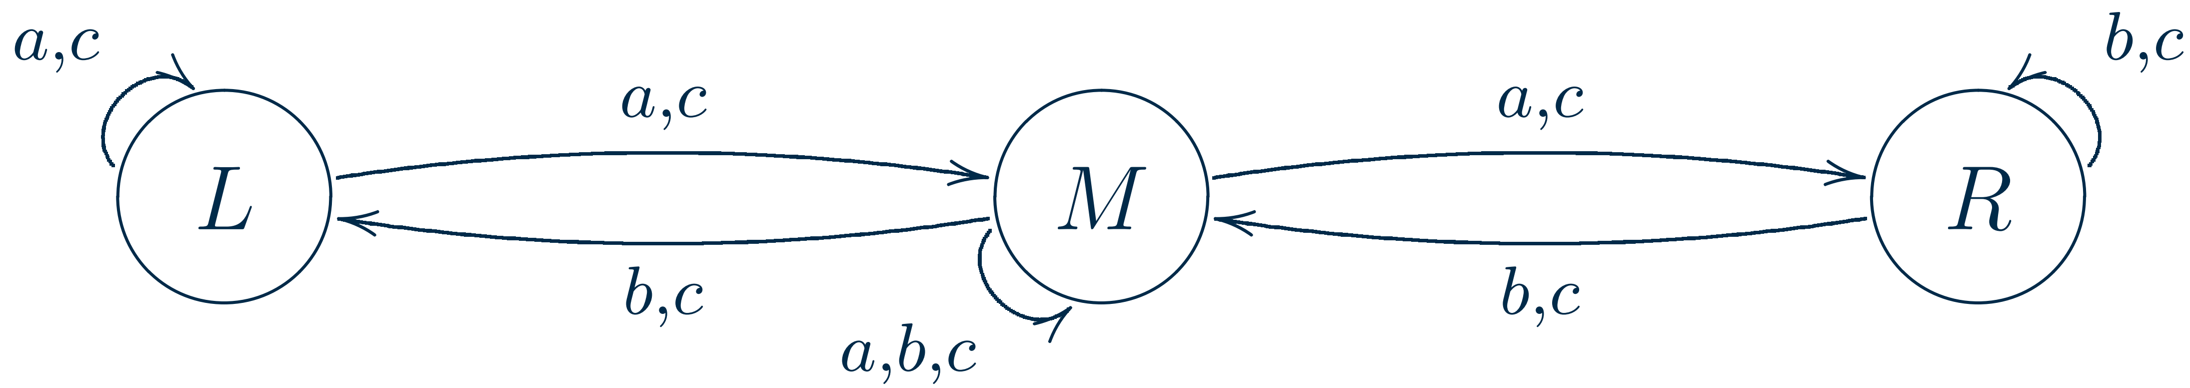

This picture is misleading: unless very special conditions hold the process cannot be compressed into an *equivalent* (?) finite-state model. In general, the transition probabilities should depend on the position.

Some remarks on this model

- This is a toy model but it exemplifies some of the issues.
- Can be used for reasoning: much better if we could have a finite-state version.
- Why not discretize right away and never worry about the continuous case?

- Because we would love the ability to refine the model later.
- A better model would be assign “rewards” or “costs” to the states and look for optimal policies rather than to think of this as a pure verification problem.
- We would benefit from more interaction with the reinforcement learning community.

Equivalence?

- What is meant by "equivalent" finite-state model?
- I assume everyone knows what is meant by ordinary bisimulation.
- We need to look into the probabilistic analogue.

Probabilistic Transition Systems

$$(S, \mathcal{A}, \forall a \in \mathcal{A} \;\; T_a : S \times S \to [0,1])$$

S: states
$\mathcal{A}$: actions
T_a: transition matrix.

This model is *reactive*: all probabilities are associated with the internal choices, no probabilities are associated with how the environment chooses the labels.

If two states behave in exactly the same way they can be combined.

Consider

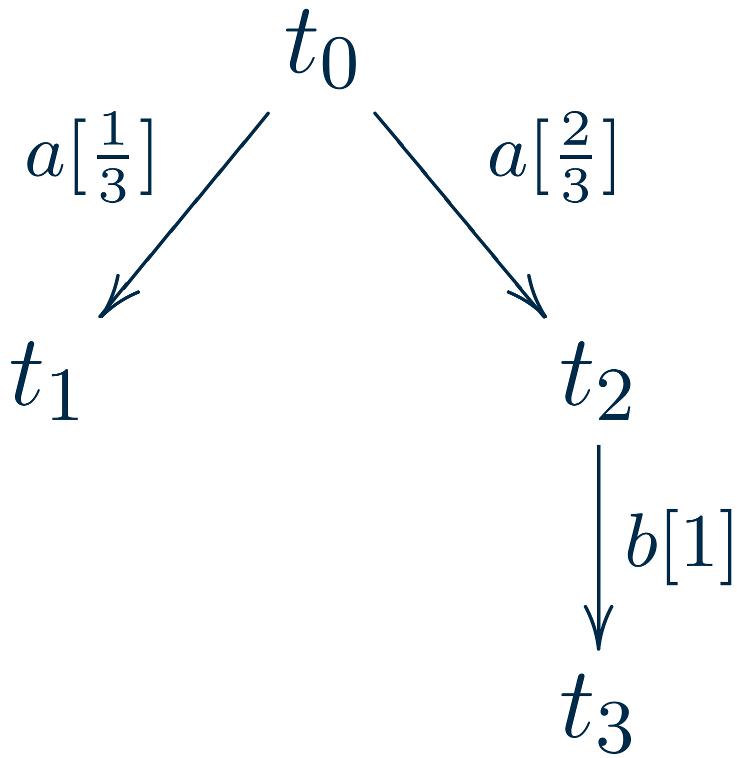

P_1

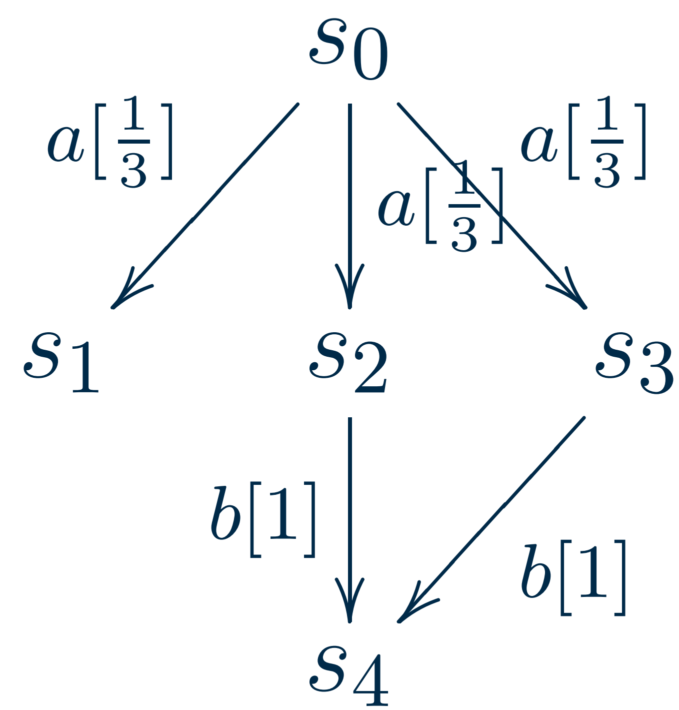

P_2

Should s_0 and t_0 be bisimilar?

Yes, but we need to add the probabilities.

Let $\mathcal{S} = (S, \mathsf{L}, T_a)$ be a PTS. An equivalence relation R on S is a **bisimulation** if whenever sRs', with $s, s' \in S$, we have that for all $a \in \mathcal{A}$ and every R-equivalence class, A, $T_a(s, A) = T_a(s', A)$.

The notation $T_a(s, A)$ means "the probability of starting from s and jumping to a state in the set A."

Two states are bisimilar if there is some bisimulation relation R relating them.

A Labelled Markov Process (LMP) is a tuple:

$$(S, \Sigma, \mathcal{A}, \forall a \in \mathcal{A} \;\; \tau_a : S \times \Sigma \rightarrow [0, 1])$$

such that

$$\tau_a(s, \cdot) \text{ is a measure and}$$

$$\tau_a(\cdot, A) \text{ is a measurable function.}$$

τ_a is just a Markov kernel.

Larsen-Skou-style bisimulation

Let $\mathcal{S} = (S, \Sigma, \mathcal{A}, \tau_a)$ be an LMP. An equivalence relation R on S is called a bisimulation if whenever sRt, then for every a and every R-closed measurable set A,

$$\tau_a(s, A) = \tau_a(t, A).$$

Two states are bisimilar if they are related by some such R.

Note that R itself does not need to satisfy any measurability condition. One only requires the matching to work on R-closed *measurable* sets.

Logical Characterization

$$\mathcal{L} ::== \mathsf{T} | \phi_1 \wedge \phi_2 | \langle a \rangle_q \phi$$

We say $s \models \langle a \rangle_q \phi$ iff

$$\exists A \in \Sigma . (\forall s' \in A . s' \models \phi) \wedge (\tau_a(s, A) > q).$$

Intuitively, we want to say $s \models \langle a \rangle_q \phi$ if $\tau_a(s, [\![\phi]\!]) > q$,
where $[\![\phi]\!] = \{s | s \models \phi\}$,
but we do not yet know that $\{s | s \models \phi\}$ is measurable.

Two systems are bisimilar iff they obey the same formulas of $\mathcal{L}$. [DEP 1998 LICS, I and C 2002]

That Cannot be Right?

$$
\begin{tikzcd}[column sep=small]
 & s_0 \arrow[dl, "a"'] \arrow[dr, "a"] & \\
s_1 & & s_2 \arrow[d, "b"] \\
 & & s_3
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
t_0 \arrow[d, "a"] \\
t_1 \arrow[d, "b"] \\
t_2
\end{tikzcd}
$$

Two processes that cannot be distinguished without negation.

The formula that distinguishes them is $\langle a\rangle(\neg\langle b\rangle\top)$.

But it is!

$$
\begin{tikzcd}
 & s_0 \arrow[dl, "{a[p]}"'] \arrow[dr, "{a[q]}"] & \\
s_1 & & s_2 \arrow[d, "b"] \\
 & & s_3
\end{tikzcd}
\qquad\qquad
\begin{tikzcd}
t_0 \arrow[d, "{a[r]}"] \\
t_1 \arrow[d, "b"] \\
t_2
\end{tikzcd}
$$

We add probabilities to the transitions.

If $p + q < r$ or $p + q > r$ we can easily distinguish them.

If $p + q = r$ and $p > 0$ then $q < r$ so $\langle a \rangle_r \langle b \rangle_1 \top$ distinguishes them.

Easy to prove that if states are bisimilar they will satisfy the same formulas (in fact in much richer logics).

For the other direction, we need to show that the relation "s and t satisfy the same formulas" is a bisimulation.

If s and t satisfy exactly the same formulas it follows that for all ϕ, $\tau_a(s, [\![\phi]\!]) = \tau_a(t, [\![\phi]\!])$.

Using some basic measure theory (Dynkin's lemma), one can show that $\tau_a(s, A) = \tau_a(t, A)$ for any A in the σ-algebra generated by the sets of the form $[\![\phi]\!]$.

At this point one gets stuck! So we make an additional assumption. The state spaces are required to be *analytic*.

Analytic spaces are extremely general: all discrete spaces, all manifolds, any complete separable metric space.

If one day you encounter a space that is not analytic you are probably just having a bad dream.

What is so great about analytic spaces and what are they anyway?

Suppose B is a measurable subset of $\mathbb{R}^2$,
its projection on the x-axis, say A, is *clearly* measurable?

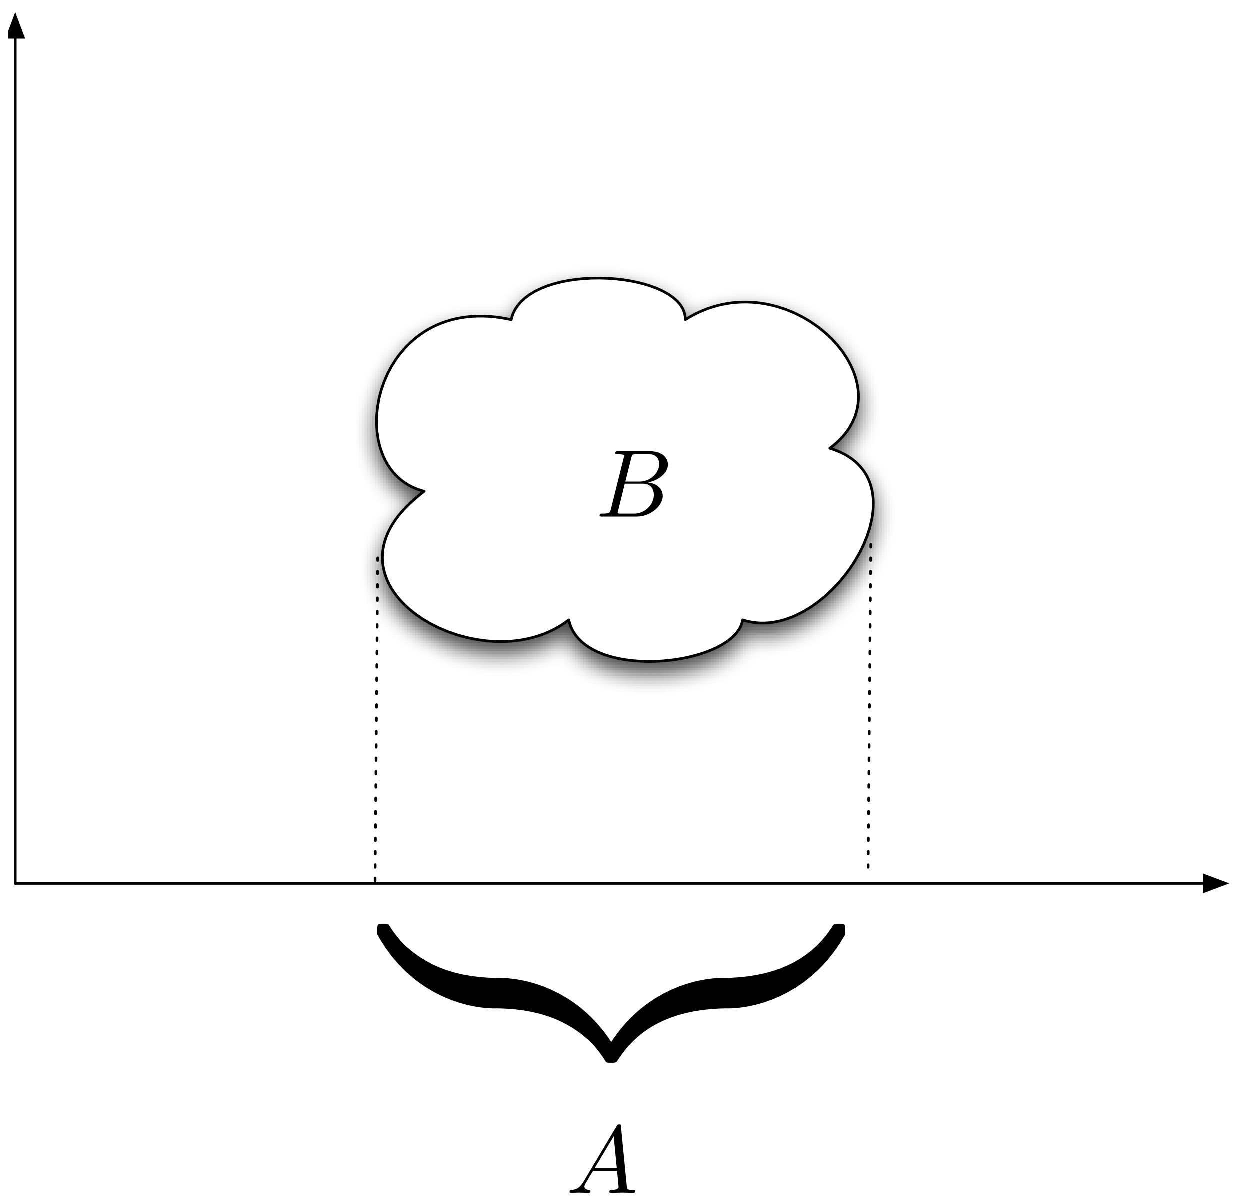

Actually, this is false! What you get is an analytic set.

Analytic sets have some magical properties.

If you take the quotient of an analytic set by a "smooth" equivalence relation you get another analytic set.

Logical equivalence is indeed a smooth equivalence relation!

Markov kernels can be defined on analytic sets.

If you have a sub-σ-algebra that is countably-generated (so not too large) and separates points (so not too small) it is the whole σ-algebra.

So the σ-algebra generated by the sets $[\![\phi]\!]$ is in fact the whole σ-algebra.

Let $\mathcal{S} = (S, \Sigma, \tau)$ be a labelled Markov process. A preorder R on S is a **simulation** if whenever sRs', we have that for all $a \in \mathcal{A}$ and every R-closed measurable set $A \in \Sigma$, $\tau_a(s, A) \leq \tau_a(s', A)$. We say s is simulated by s' if sRs' for some simulation relation R.

The logic used in the characterization has no negation, not even a limited negative construct.

One can show that if s simulates s' then s satisfies all the formulas of $\mathcal{L}$ that s' satisfies.

What about the converse?

Counter Example!

In the following picture, t satisfies all formulas of $\mathcal{L}$ that s satisfies but t does not simulate s.

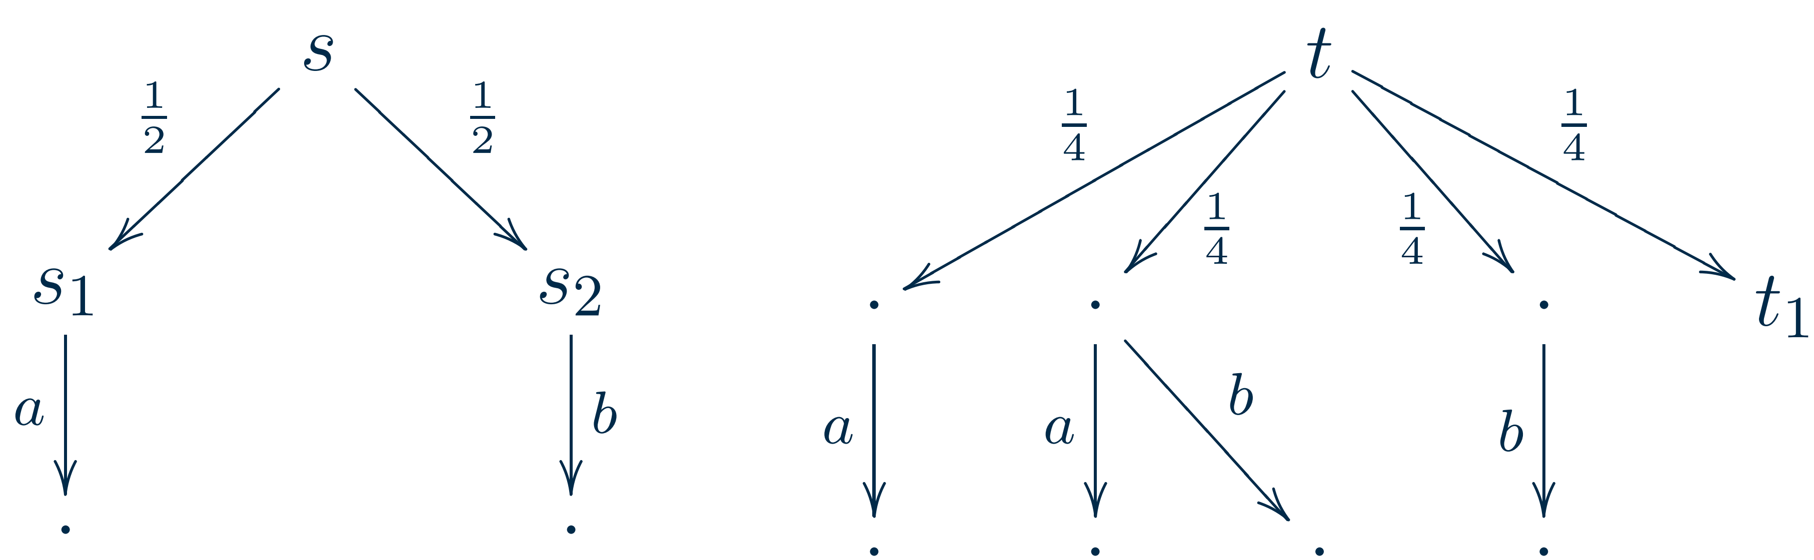

All transitions from s and t are labelled by a.

Josée Desharnais.

A formula of $\mathcal{L}$ that is satisfied by t but not by s.

$$\langle a \rangle_0 (\langle a \rangle_0 \mathsf{T} \wedge \langle b \rangle_0 \mathsf{T}).$$

A formula with disjunction that is satisfied by s but not by t:

$$\langle a \rangle_{\frac{3}{4}} (\langle a \rangle_0 \mathsf{T} \vee \langle b \rangle_0 \mathsf{T}).$$

In the following picture, t satisfies all formulas of $\mathcal{L}$ that s satisfies but t does not simulate s.

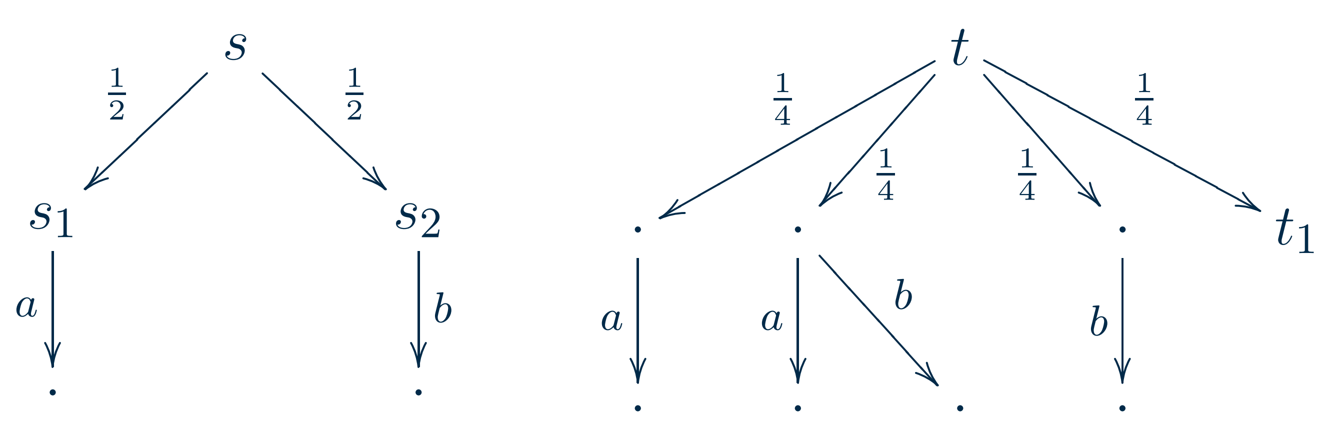

All transitions from s and t are labelled by a.

The logic $\mathcal{L}$ does **not** characterize simulation. One needs disjunction.

$$\mathcal{L}_\vee := \mathcal{L} \mid \phi_1 \vee \phi_2.$$

With this logic we have:
An **LMP** s_1 simulates s_2 if and only if for every formula ϕ of $\mathcal{L}_\vee$ we have

$$s_1 \models \phi \Rightarrow s_2 \models \phi.$$

The only proof we know uses domain theory.

Approximation Results

Our main result is a systematic approximation scheme for labeled Markov processes. The set of LMPs is a Polish space. Furthermore, our approximation results allow us to approximate integrals of continuous functions by computing them on finite approximants.

For any LMP, we explicitly provide a (countable) sequence of approximants to it such that:
- For every logical property satisfied by a process, there is an element of the chain that also satisfies the property.
- The sequence of approximants converges – in a certain metric – to the process that is being approximated.

There are lots more things to say about:

1. Metrics (Franck, Wednesday)

2. Combining probability and nondeterminism (Roberto this afternoon)

3. Approximation

4. Domains and universal Markov processes

5. Markov decision processes and applications to machine learning.

6. Real-time systems.

7. Weak bisimulation

8. Applications to information theory and security
(Pasquale later this week)

9. Duality theory

but it's lunch time!