Markov Processes as Function Transformers

Part III: Bisimulation, minimal realization and approximation

Prakash Panangaden[1]

[1]School of Computer Science
McGill University

CIRM Marseilles 20 - 24 February 2012

Outline

Today's plan

1. Defining the category **AMP**
2. Event bisimulation and zigzags
3. Bisimulation is an equivalence
4. Minimal realization
5. Logical characterization
6. Approximations
7. Projective limit and minimal realization

Key points

- We can define approximation morphisms and bisimulation morphisms in the same category.
- We can define a notion of smallest process that is bisimilar to a given process.
- We can define a notion of finite approximation and construct a projective limit of the finite approximants.
- This yields the minimal realization.

The category $\mathbf{AMP}$

- In $\mathbf{Rad}_1$ and $\mathbf{Rad}_1$ the morphisms obeyed mild conditions on the measures.
- These are sufficient to develop the functorial theory of expectation values.
- A map $\alpha : (X, p) \longrightarrow (Y, q)$ in $\mathbf{Rad}_\infty$ is said to be *measure-preserving* if $M_\alpha(p) = q$.

The category of LAMPS

We define the category $\mathbf{AMP}$ as follows. The objects consist of probability spaces (X, Σ, p), along with an abstract Markov process τ_a on X. The arrows $\alpha : (X, \Sigma, p, \tau_a) \longrightarrow (Y, \Lambda, q, \rho_a)$ are surjective measure-preserving maps from X to Y such that $\alpha(\tau_a) = \rho_a$.

- We define the category $\mathbf{Rad}_=$ to have the same objects as $\mathbf{AMP}$ but the maps are only measure preserving (and, of course, measurable).

Bisimulation traditionally

Larsen-Skou definition

Given an **LMP** (S, Σ, τ_a) an equivalence relation R on S is called a *probabilistic bisimulation* if sRt then for every *measurable* R-closed set C we have for every a

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

This variation to the continuous case is due to Josée Desharnais and her Indian friends.

Event bisimulation

- In measure theory one should focus on measurable sets rather than on *points*.
- Vincent Danos proposed the idea of *event bisimulation*, which was developed by him and Desharnais, Laviolette and P.

Event Bisimulation

Given a LMP (X, Σ, τ_a), an **event-bisimulation** is a sub-σ-algebra Λ of Σ such that (X, Λ, τ_a) is still an LMP.

- This means τ_a sends the subspace $L_\infty^+(X, \Lambda, p)$ to itself; where we are now viewing τ_a as a map on $L_\infty^+(X, \Lambda, p)$.

The bisimulation diagram

$$
\begin{array}{ccc}
L^{+}_{\infty}(X, \Sigma, p) & \xrightarrow{\tau_a} & L^{+}_{\infty}(X, \Sigma, p) \\
\uparrow\hookleftarrow & & \uparrow\hookleftarrow \\
L^{+}_{\infty}(X, \Lambda, p) & \xrightarrow{\tau_a} & L^{+}_{\infty}(X, \Lambda, p)
\end{array}
$$

Zigzag maps

We can generalize the notion of event bisimulation by using maps other than the identity map on the underlying sets. This would be a map α from (X, Σ, p) to (Y, Λ, q), equipped with LMPs τ_a and ρ_a respectively, such that the following commutes:

$$\begin{array}{ccc} L^+_\infty(X, \Sigma, p) & \xrightarrow{\tau_a} & L^+_\infty(X, \Sigma, p) \\ {\scriptstyle P_\infty(\alpha)} \uparrow & & \uparrow {\scriptstyle P_\infty(\alpha)} \\ L^+_\infty(Y, \Lambda, q) & \xrightarrow{\rho_a} & L^+_\infty(Y, \Lambda, q) \end{array} \tag{1}$$

A key diagram

When we have a zigzag the following diagram commutes:

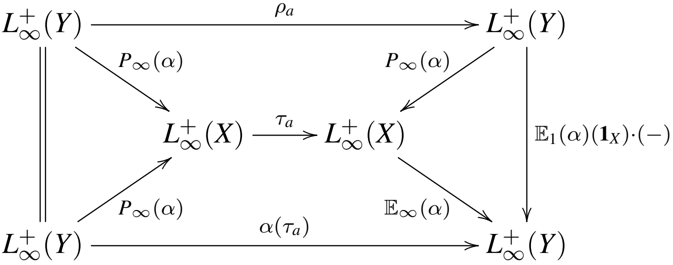

(2)

- The upper trapezium says we have a zigzag. The lower trapezium says that we have an “approximation” and the triangle on the right is an earlier lemma.
- If we “approximate” along a zigzag we actually get the exact result.
- Approximations are approximate bisimulations.

Bisimulation as a cospan

- Zigzags give a "functional" version of bisimulation; what is the relational version.
- Use co-spans of zigzags; it is usual to use spans but co-spans give a smoother and more general theory.
- With spans one can prove logical characterization of bisimulation on analytic spaces.
- With the cospan definition we get logical characterization on *all* measurable spaces.
- On analytic spaces the two concepts co-incide.
- Recent results show that the theory cannot be made to work with spans on general measurable spaces.

The official definition of bisimulation

Bisimulation

We say that two objects of $\mathbf{AMP}$, (X, Σ, p, τ) and (Y, Λ, q, ρ), are *bisimilar* if there is a third object (Z, Γ, r, π) with a pair of zigzags

$$\alpha : (X, \Sigma, p, \tau) \longrightarrow (Z, \Gamma, r, \pi)$$
$$\beta : (Y, \Lambda, q, \rho) \longrightarrow (Z, \Gamma, r, \pi)$$

giving a cospan diagram

$$\begin{array}{ccc} (X, \Sigma, p, \tau) & & (Y, \Lambda, q, \rho) \\ & {\scriptstyle\alpha}\searrow \quad \swarrow{\scriptstyle\beta} & \\ & (Z, \Gamma, r, \pi) & \end{array} \tag{3}$$

Note that the identity function on an AMP is a zigzag, and thus that any zigzag between two AMPs X and Y implies that they are bisimilar.

Fundamental categorical result

The category $\mathbf{AMP}$ has pushouts

Furthermore, if the morphisms in the span are zigzags then the morphisms in the pushout diagram are also zigzags.

Pushouts explicitly

More explicitly, let $\alpha : (X, \Sigma, p, \tau_a) \longrightarrow (Y, \Lambda, q, \rho_a)$ and $\beta : (X, \Sigma, p, \tau_a) \longrightarrow (Z, \Gamma, r, \kappa_a)$ be a span in $\mathbf{AMP}$. Then there is an object (W, Ω, μ, π_a) of $\mathbf{AMP}$ and $\mathbf{AMP}$ maps $\delta : Y \longrightarrow W$ and $\gamma : Z \longrightarrow W$ such that the diagram

$$\begin{array}{ccccc} & & (X, \Sigma, p, \tau_a) & & \\ & \swarrow^{\alpha} & & \searrow^{\beta} & \\ (Y, \Lambda, q, \rho_a) & & & & (Z, \Gamma, r, \kappa_a) \\ & \searrow_{\delta} & & \swarrow_{\gamma} & \\ & & (W, \Omega, \mu, \pi_a) & & \end{array} \tag{4}$$

commutes.

Couniversality

If (U, Ξ, ν, λ_a) is another **AMP** object and $\phi : Y \longrightarrow U$ and $\psi : Z \longrightarrow U$ are **AMP** maps such that α, β, ϕ and ψ form a commuting square, then there is a unique **AMP** map $\theta : W \longrightarrow U$ such that the diagram

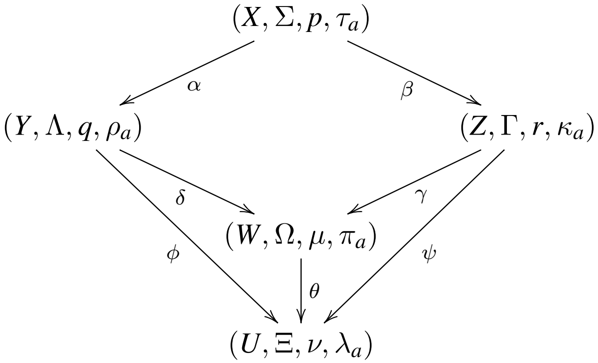

(5)

commutes.

Furthermore, if α and β are zigzags, then so are γ and δ.

Bisimulation is an equivalence

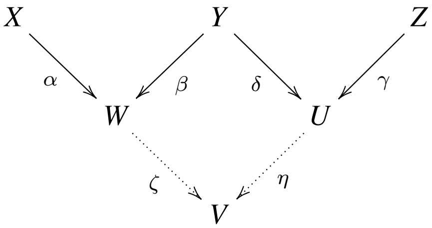

(6)

The pushouts of the zigzags β and δ yield two more zigzags ζ and η (and the pushout object V). As the composition of two zigzags is a zigzag, X and Z are bisimilar. Thus bisimulation is transitive.

What is the smallest realization of a process?

- Obviously, the concept cannot be based on counting states.
- We want to look for a bisimulation equivalent version of the process; hence with the same behaviour,
- such that any other process with the same behaviour contains this one.
- This is a classic couniversality property.

Bisimulation-minimal realization

Definition of minimal realization

Given an AMP (X, Σ, p, τ_a), a **bisimulation-minimal realization** of this abstract Markov process is an AMP $(\tilde{X}, \Gamma, r, \pi_a)$ and a zigzag in $\mathbf{AMP}$ $\eta : X \longrightarrow \tilde{X}$ such that for every zigzag β from X to another AMP (Y, Λ, q, ρ_a), there is a zigzag γ from (Y, Λ, q, ρ_a) to $(\tilde{X}, \Gamma, r, \pi_a)$ with $\eta = \gamma \circ \beta$.

If we think of a zigzag as defining a quotient of the original space then $\tilde{X}$ is the “most collapsed” version of X.

Existence theorem

Given any AMP (X, Σ, p, τ_a) there exists another AMP $(\tilde{X}, \Gamma, r, \pi_a)$ and a zigzag η in $\mathbf{AMP}$, $\eta : X \longrightarrow \tilde{X}$ such that $(\tilde{X}, \Gamma, r, \pi_a)$ and η define a bisimulation-minimal realization of (X, Σ, p, τ_a).

Proof idea: Intersect all event bisimulations to get a smallest (fewest sets in the σ-algebra) event bisimulation. Define the associated equivalence relation and form the quotient.

Two AMPs and are bisimilar if and only if their minimal realizations and respectively are isomorphic.

A modal logic

We define a logic $\mathcal{L}$ as follows, with $a \in \mathcal{A}$:

$$\mathcal{L} ::= \mathbf{T} | \phi \wedge \psi | \langle a \rangle_q \psi$$

Given a labelled AMP (X, Σ, p, τ_a), we associate to each formula ϕ a measurable set $[\![\phi]\!]$, defined recursively as follows:

$$\begin{aligned} &[\![\mathbf{T}]\!] = X \\ &[\![\phi \wedge \psi]\!] = [\![\phi]\!] \cap [\![\psi]\!] \\ &\left[\!\!\left[\langle a \rangle_q \psi\right]\!\!\right] = \left\{s : \tau_a(\mathbf{1}_{[\![\psi]\!]})(s) > q\right\} \end{aligned}$$

We let $[\![\mathcal{L}]\!]$ denotes the measurable sets obtained by all formulas of $\mathcal{L}$.

Logical characterization of bisimulation

Main theorem

Given a LAMP (X, Σ, p, τ_a), the σ-field $\sigma([\![\mathcal{L}]\!])$ generated by the logic $\mathcal{L}$ is the smallest event-bisimulation on X. That is, the map $i : (X, \Sigma, p, \tau_a) \longrightarrow (X, \sigma([\![\mathcal{L}]\!]), p, \tau_a)$ is a zigzag; furthermore, given any zigzag $\alpha : (X, \Sigma, p, \tau_a) \longrightarrow (Y, \Lambda, q, \rho_a)$, we have that $\sigma([\![\mathcal{L}]\!]) \subseteq \alpha^{-1}(\Lambda)$.

Hence, the σ-field obtained on X by the smallest event bisimulation is precisely the σ-field we obtain from the logic.

The approximation map (recap)

The expectation value functors project a probability space onto another one with a possibly coarser σ-algebra.
Given an AMP on (X,p) and a map $\alpha : (X,p) \to (Y,q)$ in $\mathbf{Rad}_\infty$, we have the following approximation scheme:

Approximation scheme

$$\begin{array}{ccc} L^+_\infty(X,p) & \xrightarrow{\tau_a} & L^+_\infty(X,p) \\ {\scriptstyle P_\infty(\alpha)}\uparrow & & \downarrow{\scriptstyle \mathbb{E}_\infty(\alpha)} \\ L^+_\infty(Y,q) & \xrightarrow{\alpha(\tau_a)} & L^+_\infty(Y,q) \end{array}$$

A special case (recap)

- Take (X, Σ) and (X, Λ) with $\lambda \subset \Sigma$ and use the measurable function $id : (X, \Sigma) \longrightarrow (X, \Lambda)$ as α.

Coarsening the σ-algebra

$$\begin{array}{ccc} L^+_\infty(X, \Sigma, p) & \xrightarrow{\tau_a} & L^+_\infty(X, \Sigma, p) \\ {\scriptstyle P_\infty(\alpha)}\uparrow & & \downarrow{\scriptstyle \mathbb{E}_\infty(\alpha)} \\ L^+_\infty(X, \Lambda, p) & \xrightarrow{id(\tau_a)} & L^+_\infty(X, \Lambda, p) \end{array}$$

- Thus $id(\tau_a)$ is the approximation of τ_a obtained by averaging over the sets of the coarser σ-algebra Λ.

Finite approximations

Let (X, Σ, p, τ_a) be a LAMP. Let $\mathcal{P} = 0 < q_1 < q_2 < \ldots < q_n < 1$ be a finite partition of the unit interval with each q_i a rational number. We call these *rational partitions*. We define a family of finite π-systems, subsets of Σ, as follows:

$$\begin{aligned}
\Phi_{\mathcal{P},0} &= \{X, \emptyset\} \\
\Phi_{\mathcal{P},n} &= \pi\left(\left\{\tau_a(\mathbf{1}_A)^{-1}(q_i, 1] : q_i \in \mathcal{P}, A \in \Phi_{\mathcal{P},n-1}, a \in \mathcal{A}\right\} \cup \Phi_{\mathcal{P},n-1}\right) \\
&= \pi\left(\left\{\left[\!\left[\langle a \rangle_{q_i} \mathbf{1}_A\right]\!\right] : q_i \in \mathcal{P}, A \in \Phi_{\mathcal{P},n-1}, a \in \mathcal{A}\right\} \cup \Phi_{\mathcal{P},n-1}\right)
\end{aligned}$$

where $\pi(\Omega)$ means the π-system generated by the family of sets Ω.

Approximation pairs

For each pair $(\mathcal{P}, M)$ consisting of a rational partition and a natural number, we define a σ-algebra $\Lambda_{\mathcal{P},M}$ on X as $\Lambda_{\mathcal{P},M} = \sigma(\Phi_{\mathcal{P},M})$, the σ-algebra generated by $\Phi_{\mathcal{P},M}$. We call each pair $(\mathcal{P}, M)$ consisting of a rational partition and a natural number an *approximation pair*.

Approximation and Logic

The following result links the finite approximation with the formulas of the logic used in the characterization of bisimulation.

Crucial fact

Given any labelled AMP (X, Σ, p, τ_a), the σ-algebra $\sigma\left(\bigcup \Phi_{\mathcal{P},M}\right)$, where the union is taken over all approximation pairs, is precisely the σ-algebra $\sigma\left[\!\left[\mathcal{L}\right]\!\right]$ obtained from the logic.

Relating finite approximations

- Given two approximation pairs such that $(\mathcal{P}, M) \leq (\mathcal{Q}, N)$, we have a map
- $$i_{(\mathcal{Q},N),(\mathcal{P},M)} : (X, \Lambda_{\mathcal{Q},N}, \Lambda_{\mathcal{Q},N}(\tau_a)) \longrightarrow (X, \Lambda_{\mathcal{P},M}, \Lambda_{\mathcal{P},M}(\tau_a))$$
- which is well defined by the inclusion $\Lambda_{\mathcal{P},M} \subseteq \Lambda_{\mathcal{Q},N} \subseteq \Sigma$.
- Furthermore if $(\mathcal{P}, M) \leq (\mathcal{Q}, N) \leq (\mathcal{R}, K)$ the maps compose to give
$$i_{(\mathcal{R},K),(\mathcal{P},M)} = i_{(\mathcal{R},K),(\mathcal{Q},N)} \circ i_{(\mathcal{Q},N),(\mathcal{P},M)}.$$
- In short we have a projective system of such maps indexed by our poset of approximation pairs.

Finite spaces

- We define the space $\hat{X}_{\mathcal{Q},N}$ as the quotient of X by the equivalence relation that identifies two points that cannot be separated by measurable sets of $\Lambda_{\mathcal{Q},N}$.
- These spaces have finitely many points.
- The quotient map $q : X \longrightarrow \hat{X}_{\mathcal{Q},N}$ induces a projected version of the LAMP τ_a.
- When the approximations are refined the quotients compose so we can define maps between quotient spaces.

Projective diagram fragment

We get the following commuting diagram:

$$\begin{array}{ccc}
(X, \Lambda_{\mathcal{Q},N}, \Lambda_{\mathcal{Q},N}(\tau_a)) & \xrightarrow{i_{(\mathcal{Q},N),(\mathcal{P},M)}} & (X, \Lambda_{\mathcal{P},M}, \Lambda_{\mathcal{P},M}(\tau_a)) \\
\downarrow \pi_{\mathcal{Q},N} & & \downarrow \pi_{\mathcal{P},M} \\
(\hat{X}_{\mathcal{Q},N}, \phi_{\mathcal{Q},N}(\tau_a)) & \xrightarrow[j_{(\mathcal{Q},N),(\mathcal{P},M)}]{} & (\hat{X}_{\mathcal{P},M}, \phi_{\mathcal{P},M}(\tau_a))
\end{array} \tag{7}$$

Existence of a projective limit

Main theorem

The probability spaces of finite approximants $\hat{X}_{\mathcal{P},M}$ of a measure space (X, Σ, p, τ_a) each equipped with the discrete σ-algebra (i.e. the σ-algebra of all subsets) indexed by the approximation pairs, form a projective system in the category $\mathbf{Rad}_=$. This system of finite approximants to the LAMP (X, Σ, p, τ_a) has a projective limit in the category $\mathbf{Rad}_=$.

This uses a theorem of Choksi from 1958. In typical analysis style, he constructs the required limit but does not prove any universal property. It was a non-trivial extension to show this.

Picture of the situation

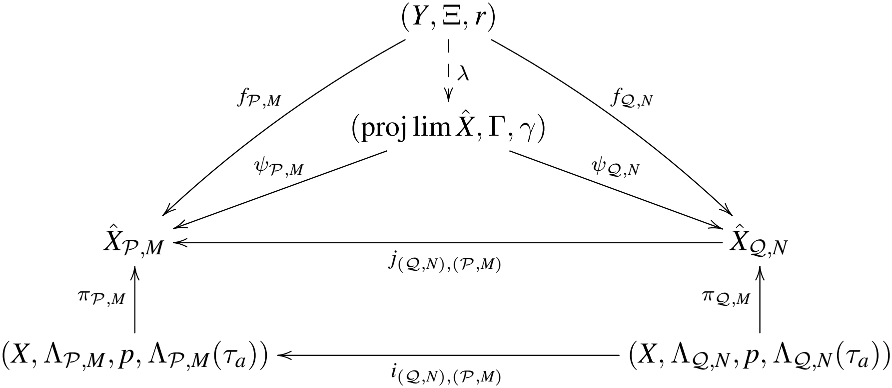

What can we say about the LAMP?

We can now consider the LAMP structure. We do not get a universal property in the category $\mathbf{AMP}$, however, the universality of the construction in $\mathbf{Rad}_=$ almost forces the structure of a LAMP on the projective limit constructed in $\mathbf{Rad}_=$.

LAMP on the projective limit

A LAMP can be defined on the projective limit constructed in $\mathbf{Rad}_=$ so that the cone formed by this limit object and the maps to the finite approximants yields a commuting diagram in the category $\mathbf{AMP}$.

Approximation and minimal realization

- The LAMP obtained by forming the projective limit in the category $\mathbf{Rad}_=$ and then defining a LAMP on it is isomorphic to the minimal realization of the original LAMP.
- This gives a very pleasing connection between the approximation process and the minimal realization.

Two routes to the minimal realization

Given an AMP (X, Σ, p, τ_a), the projective limit of its finite approximants $(\text{proj} \lim \hat{X}, \Gamma, \gamma, \zeta_a)$ is isomorphic to its minimal realization $(\tilde{X}, \Xi, r, \xi_a)$.

Summary

- Viewing Markov processes as function transformers
- The old theory can be redone more smoothly and with better results
- Approximation via averaging makes sense in theory and practice

The future

- A general theory with all L_p spaces.
- Tying up with Stone duality; much work in progress.
- Projective limit in $\mathbf{AMP}$?
- Continuous time?