

Fixed Points for Markov Decision Processes

Johannes Hölzl *

TU München
hoelzl@in.tum.de

Abstract

The goal of this paper is to advertise the application of fixed points and ω -complete partial orders (ω -cpo) in the formalization and analysis of probabilistic programming languages. The presented work is formalized in the Isabelle theorem prover.

By applying ω -cpo to the analysis of MDPs we get a nice theory of fixed points. This allows us to transfer least and greatest fixed points through expectation on Markov chains and maximal and minimal expectation on MDPs. One application is to define the operational semantics of pGCL by MDPs, e.g. relating the denotational and operational semantics of pGCL is now a matter of fixed point equations and induction.

1. Introduction

It is standard to use fixed points to describe solutions for probabilistic model checking [2], or to give semantics for probabilistic programming languages as fixed points [5]. This is similar to the work by Monniaux [7]. Both are fixed points on complete lattices (either $[0, 1]$ or $\overline{\mathbb{R}}^+$, the non-negative reals with ∞). We are concerned with a more general class, namely ω -complete partial orders (ω -cpo). Here fixed points are only well-defined for continuous functions. A function F is sup-continuous (and inf-continuous) if for all increasing (and decreasing) sequences X , the function F commutes with supremum (and infimum) of X :

$$F(\bigsqcup_i X_i) = \bigsqcup_i F(X_i) \quad \text{and} \quad F(\prod_i X_i) = \prod_i F(X_i). \quad (1)$$

The least (and greatest) fixed point of F is equal to the supremum (and infimum) of the iterations of F , when F is continuous:

$$\text{lfp } F = \bigsqcup_i F^i(\perp) \quad \text{gfp } F = \prod_i F^i(\top) \quad (2)$$

Hence for continuous functions F, G, α on ω -cpo we derive a *transfer rule* similar to the fusion rule in [4]:

$$\frac{\alpha(\perp) = \perp \quad \alpha \circ G = F \circ \alpha}{\alpha(\text{lfp } G) = \text{lfp } F} \quad (3)$$

We get a dual rule for greatest fixed points.

We need ω -cpo as they are compatible with measure theory: the Borel-measurable functions $\mathcal{M}(\mathcal{A}, \mathcal{B}(\overline{\mathbb{R}}^+))$ are not a complete lattice, but a ω -cpo, for each measurable spaces \mathcal{A} . We know that least (and greatest) fixed points $\text{lfp } F$ (and $\text{gfp } F$) are measurable assuming the functional F is continuous and measurable:

$$F \in \mathcal{M}(\mathcal{A}, \mathcal{B}(\overline{\mathbb{R}}^+)) \rightarrow \mathcal{M}(\mathcal{A}, \mathcal{B}(\overline{\mathbb{R}}^+)) \quad (4)$$

In the rest of this paper we always assume that functions are measurable. From Lebesgue's monotone convergence theorem follows that Lebesgue integration is a continuous function. There is a similar monotone convergence theorem for decreasing sequences. However, this rules requires that the integral is finite starting at some

point i in the sequence X_i . We derive the following transfer rule on the Lebesgue integral:

$$\frac{F, G \text{ sup-continuous} \quad \forall s, f \leq \text{lfp } F. \int F f dM_s = G(\lambda s. \int f dM_s)}{\int \text{lfp } F dM_s = \text{lfp } G s} \quad (5)$$

Compared to Rule (3) we provide it with a stronger equation (restricted to $f \leq \text{lfp } F$), as this is necessary for certain proofs. Also adding the index s to the family of probability measures M_s will be very handy when working on the trace spaces of Markov chains and MDPs. We also assume measurability of F as the one in Rule (4). The rule for the greatest fixed point is very similar, however it requires that the integral is always finite.

$$\frac{F, G \text{ inf-continuous} \quad \forall f, s. \int F f dM_s < \infty \quad \forall s, f \geq \text{gfp } F. \int F f dM_s = G(\lambda s. \int f dM_s)}{\int \text{gfp } F dM_s = \text{gfp } G s} \quad (6)$$

When solving integrals of fixed points we can now apply Rule (5) or Rule (6) and then continue with fixed point equations like rolling, diagonalisation, iteration, or induction [4]. In the following sections we apply these rules on Markov chains, MDPs and for the semantics of pGCL.

2. Markov chains

A Markov chain is given by its Markov kernel $K : A \rightarrow \mathcal{D}(A)$, where $\mathcal{D}(A)$ is a discrete probability measure on A . Given a kernel K we construct the trace space T_s (e.g. by the Ionescu-Tulcea extension theorem [8] or similar) starting in state s with the iteration property:

$$\int f dT_s = \int_t (\int_\omega f(t\omega) dT_t) dK_s \quad (7)$$

Here $t\omega$ prepends t to the sequence ω . Using least and greatest fixed points we define the usual LTL operators, extending the work from Blanchette et al. [3], but also first hitting time or cumulated rewards. For example eventually \diamond and first hitting time h are defined as follows:

$$\begin{aligned} \diamond(X) &= \text{lfp } (\lambda F. \lambda(t\omega). \text{if } t \in X \text{ then } 1 \text{ else } F(\omega)) \\ h(X) &= \text{lfp } (\lambda F. \lambda(t\omega). \text{if } t \in X \text{ then } 0 \text{ else } 1 + F(\omega)) \end{aligned}$$

As the functionals in the definitions are continuous and measurable, it follows that \diamond and h are also measurable functions. With the transfer rule (3) we get the linear equation system known from probabilistic model checking [2].

The transfer rules are also used in proofs, for example to prove fairness (originally from [1]) we show that the not fairness set is expressed as gfp , and bound from above by 0. Similarly, we prove that hitting time $h(X)$ is finite, assuming a finite Markov chain, and that the set X is reachable.

* Supported by the DFG Projekt NI 491/15-1

3. Markov decision processes

A Markov decision process (MDP) is similarly given as a kernel $K : A \rightarrow \mathcal{P}(\mathcal{D}(A))$. This time the trace space is parameterized by a scheduler $\sigma(s)$ starting in s .

$$\int f dT_{\sigma(s)} = \int_t \left(\int_{\omega} f(t \cdot \omega) dT_{\sigma(s,t)} \right) dK_{\sigma(s)} \quad (8)$$

Here $\sigma(s, t)$ is the continuation of the scheduler $\sigma(s)$ going to t , and $K_{\sigma(s)}$ is a distribution in K_s chosen by the scheduler $\sigma(s)$. Now the minimal and maximal expectation is defined as:

$$E_s^{\min}(f) = \prod_{\sigma(s)} \int f dT_{\sigma(s)} \quad E_s^{\max}(f) = \bigsqcup_{\sigma(s)} \int f dT_{\sigma(s)} \quad (9)$$

For these we get also iterative rules similar to Rule (7):

$$E_s^{\max}(f) = \bigsqcup_{\mu \in K_s} \int_t E_t^{\max}(\omega. f(t \cdot \omega)) d\mu \quad (10)$$

For E^{\min} we get transfer rule for lfp, and for E^{\max} we get the transfer rule only when the sets K_s are finite for all s . From this again we get directly the fixed point we can use to get the linear equation system used for model checking.

4. pGCL Semantics

The probabilistic guarded command language (pGCL) is a imperative language with probabilistic and non-deterministic choice. Gretz et al. [5] compare the usual (denotational) expectation transformer semantics with operational semantics. As operational semantics they define for each pGCL program a MDP. Then finally they show that the reward of the MDP equals the expectation of the expectation transformer. The proof works by representing the expected reward as the sum over all possible paths. In this section we want to show an alternative proof, representing the reward as a fixed point. First we present the pGCL semantics similar to [5]:

$$\begin{aligned} pgcl &:= Skip \mid Abort \mid Assign(\sigma \rightarrow \sigma) \\ &\mid Seq\ pgcl\ pgcl \mid Par\ pgcl\ pgcl \\ &\mid If(\sigma \rightarrow bool)\ pgcl\ pgcl \\ &\mid Prob[0, 1]\ pgcl\ pgcl \\ &\mid While(\sigma \rightarrow bool)\ pgcl\ pgcl \end{aligned}$$

The weakest pre-expectation $wp\ c\ f$ is an expectation transformer recursively computed by the program structure of c :

$$\begin{aligned} wp : pgcl &\rightarrow (\sigma \rightarrow \mathbb{R}^+) \rightarrow (\sigma \rightarrow \mathbb{R}^+) \\ wp\ Skip\ f &= f \\ wp\ Abort\ f &= \perp \\ wp\ (Assign\ u)\ f &= f \circ u \\ wp\ (Seq\ c_1\ c_2)\ f &= wp\ c_1\ (wp\ c_2\ f) \\ wp\ (Par\ c_1\ c_2)\ f &= wp\ c_1\ f \sqcap wp\ c_2\ f \\ wp\ (If\ b\ c_1\ c_2)\ f &= \\ &\lambda s. \text{if } b\ s \text{ then } wp\ c_1\ f\ s \text{ else } wp\ c_2\ f\ s \\ wp\ (Prob\ p\ c_1\ c_2)\ f &= \\ &\lambda s. p \cdot wp\ c_1\ f\ s + (1-p) \cdot wp\ c_2\ f\ s \\ wp\ (While\ b\ c)\ f &= \\ &\text{lfp } (\lambda X\ s. \text{if } b\ s \text{ then } wp\ c\ X\ s \text{ else } f\ s) \end{aligned}$$

The operational semantics are K used to construct the MDP:

$$\begin{aligned} K : (pgcl \times \sigma) &\rightarrow \mathcal{P}(\mathcal{D}(pgcl \times \sigma)) \\ K\ (Skip, s) &= \ll Skip, s \gg \\ K\ (Abort, s) &= \ll Abort, s \gg \\ K\ (Assign\ u, s) &= \ll Skip, u\ s \gg \\ K\ (Seq\ c_1\ c_2, s) &= \\ &[K(c_1, s)](\lambda(c'_1, s'). \text{if } c'_1 = Skip \text{ then } c_2 \text{ else } Seq\ c'_1\ c_2, s') \\ K\ (Par\ c_1\ c_2, s) &= \ll c_1, s \gg \cup \ll c_2, s \gg \\ K\ (If\ b\ c_1\ c_2, s) &= \text{if } b\ s \text{ then } K(c_1, s) \text{ else } K(c_2, s) \\ K\ (Prob\ p\ c_1\ c_2, s) &= \{\{p \mapsto (c_1, s), (1-p) \mapsto (c_2, s)\}\} \\ K\ (While\ b\ c, s) &= \\ &\text{if } b\ s \text{ then } \ll Seq\ c\ (While\ b\ c), s \gg \text{ else } \ll Skip, s \gg \end{aligned}$$

Here $\ll c, s \gg$ is the singleton set, containing the Dirac distribution over (c, s) . The syntax $[K]f$ maps the elements of K over f , i.e. if K has type $\mathcal{P}(\mathcal{D}(A))$ and $f : A \rightarrow B$, then the result type is $\mathcal{P}(\mathcal{D}(B))$.

Next we define the reward of f on a trace as a least fixed point:

$$r\ f = \text{lfp } (\lambda F\ ((c, s) \cdot \omega). \text{if } c = Skip \text{ then } f(s) \text{ else } F(\omega))$$

Using the transfer rule on E^{\min} we get the following equation:

$$\begin{aligned} E_x^{\min}(r\ f) &= \text{lfp } (\lambda F\ x. \\ &\prod_{\mu \in K_x} \int_{(c,s)} \text{if } c = Skip \text{ then } f(s) \text{ else } F(c, s) d\mu) x \end{aligned} \quad (11)$$

With this we finally prove:

$$E_{(c,s)}^{\min}(r\ f) = wp\ c\ f\ s \quad (12)$$

Before we start the proof by induction, we prove the *Seq*-case:

$$E_{(Seq\ c_1\ c_2, s)}^{\min}(r\ f) = E_{(c_1, s)}^{\min} \left(r \left(\lambda s'. E_{(c_2, s')}^{\min}(r\ f) \right) \right)$$

This is proved by rewriting with Eq. (11) and then applying fixed point induction in both directions.

Now we prove Eq. (12), by induction over c and generalizing in s and f . The only interesting case is the (While $g\ c$)-case. That wp is less or equal E^{\min} follows from the fact that E^{\min} fullfills already the fixed point equation. For the other direction we unroll the while-loop to the following functional:

$$\begin{aligned} w\ F\ s &= \prod_{\mu \in K_x} \int_{(d,s)} \text{if } d = Skip \text{ then if } g\ s \text{ then } F(c, s) \\ &\text{else } f(s) \text{ else } F(d, s) d\mu \end{aligned}$$

We finish with $E_{(While\ g\ c, s)}^{\max}(r\ f) \leq \text{lfp } w(c, s) = wp(While\ g\ c)\ s$, where first part is proved by induction and the second by loop unrolling and the diagonal rule for least fixed points.

5. Formalization in Isabelle

The theory presented in this paper is formalized in Isabelle, and can be found in the AFP entry `Markov_Models` [6]. Compared to the presentation in this paper, many statements in Isabelle are more complicated, as we need to explicitly state measurability. Otherwise, the fixed point theorems on Markov chains and MDPs simplified the formalization considerable as many ε -proofs are now reduced to induction on fixed points, or even direct rewriting by fixed point equations. Proofs like fairness on Markov chains or finite hitting time on finite Markov chains were quite technical before and got more concise by using fixed point equations.

References

- [1] C. Baier. *On the Algorithmic Verification of Probabilistic Systems*. Habilitation, Universität Mannheim, 1998.
- [2] C. Baier and J.-P. Katoen. *Principles of Model Checking*. The MIT Press, Cambridge, Massachusetts, 2008.
- [3] J. C. Blanchette, A. Popescu, and D. Traytel. Unified classical logic completeness. In S. Demri, D. Kapur, and C. Weidenbach, editors, *Automated Reasoning*, volume 8562 of *LNCs*, pages 46–60. 2014.
- [4] B. A. Davey and H. A. Priestley. *Introduction to Lattices and Order*. Cambridge University Press, second edition, 2002.
- [5] F. Gretz, J. Katoen, and A. McIver. Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Performance Evaluation*, 73:110–132, 2014.
- [6] J. Hölzl and T. Nipkow. Markov models. *The Archive of Formal Proofs*, Jan 2012. http://afp.sf.net/entries/Markov_Models.shtml.
- [7] D. Monniaux. Abstract interpretation of programs as Markov decision processes. In *SAS 2003*, volume 58 of *Science of Computer Programming*, pages 179–205. 2005.
- [8] D. Pollard. *A User's Guide to Measure Theoretic Probability*. Cambridge Series in Statistical and Probabilistic Mathematics. 2002.