# The semantics of the probabilistic Maude strategy language

Technical report (draft) — September, 2022

Rubén Rubio, Narciso Martí-Oliet, Isabel Pita, Alberto Verdejo
Facultad de Informática, Universidad Complutense de Madrid, Spain
{rubenrub,narciso,ipandreu,jalberto}@ucm.es

This document describes the semantics of the probabilistic Maude strategy language as an extension of the semantics of the standard Maude strategy language in [4]. The probabilistic version of the Maude strategy language includes all operators of the original language and the following new three:

- `choice(`$w_1$ `:` $\alpha_1$`,` ..., $w_n$ `:` $\alpha_n$`)` that selects one of the strategies $\alpha_k$ according to their weights $w_k$. These weights are terms in the `Nat` or `Float` sorts that may contain variables if they are defined in the outer scope. This is an evolution of the nondeterministic choice operator $\alpha_1$ `|` $\cdots$ `|` $\alpha_n$, and similar constructs exist in a probabilistic extension of ELAN [3] and in Porgy [1].

- `sample` $X$ `:=` $\pi$`(`$t_1$`,` ..., $t_n$`)` `in` $\alpha$ that samples the variable $X$ from a probabilistic distribution with parameters $t_1, \ldots, t_n$ that may contain variables defined in the outer contexts. The new variable $X$ can be freely used in $\alpha$. Both the variable $X$ and the parameters must be of sort `Float`. The available distributions are `bernouilli(`$p$`)`, `uniform(`$a$`,` $b$`)`, `exp(`$\lambda$`)`, `norm(`$\mu$`,` $\sigma$`)`, and `gamma(`$\alpha$`,` $\beta$`)`.

- An extension of the `matchrew`, `xmatchrew`, and `amatchrew` combinators of the standard strategy language to allow specifying the weight of every match and selecting one according to these weights. Syntactically, an optional infix `with weight` $w$ is added to the original operators, like in

  `matchrew` $P(X_1, \ldots, X_n)$ `s.t.` $C$ `with weight` $w$ `by` $X_1$ `using` $\alpha_1$`,` ..., $X_n$ `using` $\alpha_n$,

  where the weight $w$ is a term of sort `Nat` or `Float` that may contain variables from the matching, the condition, and the outer scope.

## 1 Small-step operational semantics

The small-step operational semantics of the Maude strategy language in [4] is defined on top of *execution states* $\mathcal{XS}$ holding both the term being rewritten and the strategy continuation. Steps are partioned into *control steps* $\to_c$ that advance the execution of the strategy without modifying the current term, and *system steps* $\to_s$ that apply a single rule rewrite on the current term. The rules of the operational semantics define these two small-step relations, but the most meaningful relation is the derived $\twoheadrightarrow = \to_s \circ \to_c^*$ that consists of a single system step preceded by as many control transitions as needed. Since the strategy language admits *delayed failures*, i.e. discarding an unbounded number of previous steps with the `fail` strategy or other implicit failures, not all steps described by the small-step operational semantics are valid. The validity of a state and consequently of a transition leading to it is defined as

$$\text{valid}(q) = (\exists\, t \in T_\Sigma \quad q \to_{s,c}^* t \,@\, \varepsilon) \;\vee\; (\exists\, x : \mathbb{N} \to \mathcal{XS} \quad x(0) = q \wedge \forall n \in \mathbb{N} \quad x(n) \twoheadrightarrow x(n+1))$$

where $\to_{s,c} = \to_s \cup \to_c$. States of the form $t \,@\, \varepsilon$ and the corresponding terms $t$ are called *solutions* of the strategy. Thus, an execution state is valid if it leads to a solution or to a nonterminating rewriting path.

For extending the small-step operational semantics to the new probabilistic operators, we first provide *nondeterministic* rules for all of them.

The `choice` operator may execute any of its substrategies $\alpha_k$ whose weights evaluate to a positive number:

$$t \,@\, \texttt{choice(}w_1 \texttt{ : } \alpha_1\texttt{, } \ldots\texttt{, } w_n \texttt{ : } \alpha_n\texttt{)}\, s \to_c t \,@\, \alpha_k\, s \qquad \text{for } 1 \leq k \leq n \text{ and } \theta(w_k) > 0$$

where $\theta = \mathrm{vctx}(s)$ denotes the current variable context implied by $s$. The precise meaning of $\theta(w_k) > 0$ is that the term $\theta(w_k)$ is equivalent modulo equations either to a non-zero constant of sort `Nat` or to a positive constant of sort `Float`.

For the `sample` operator, the variable $X$ can take any value in the domain of the probabilistic distribution for the given parameters:

$$(\texttt{sample } X \texttt{ := } \pi\texttt{(}t_1\texttt{, } \ldots\texttt{, } t_n\texttt{) in } \alpha)\, s \to_c t \,@\, \alpha\, \theta[X/x]\, s \qquad \text{for } x \in \operatorname{dom} \pi(\theta(t_1), \ldots, \theta(t_n))$$

Notice that there may be infinitely many, even uncountably many successors from a `sample` state. For convenience, we consider that this rule is only applicable when the measure of the set of valid successors of the state is positive, $\nu(\{x \in \mathbb{R} : \mathrm{valid}(t \,@\, \alpha\, \theta[X/x]\, s)\}) > 0$.

Finally, the semantic rule for the `matchrew with weight` operator coincides with that of the standard `matchrew` [4] with the additional condition that the weight $\theta(w)$ must be positive:

$$\begin{aligned} & t \,@\, (\texttt{matchrew } P \texttt{ s.t } C \texttt{ with weight } w \texttt{ by } x_1 \texttt{ using } \alpha_1, \ldots, x_n \texttt{ using } \alpha_n)\, s \\ & \quad \to_c \mathrm{subterm}(x_1 : \sigma(x_1) \,@\, \alpha_1\, \sigma, \ldots, x_n : \sigma(x_n) \,@\, \alpha_n\, \sigma; \sigma_{-\{x_1,\ldots,x_n\}}(P)) \,@\, s \\ & \qquad \text{if } \sigma \in \mathsf{mcheck}(P, t, C, \theta) \text{ and } \sigma(w) > 0 \end{aligned}$$

Similar rules are valid for the other variants, `xmatchrew` and `amatchrew`.

None of the previous rules talks about probabilities, but we will turn the semantic graph into a probabilistic transition system with the following definitions.

**Definition 1** (PARS)**.** A *probabilistic abstract reduction system* (PARS) [2] is a tuple $(S, R)$ where $S$ is a measurable space of states and $R \subseteq S \times \mathcal{M}(S)$ is a transition relation that associates probability measures over $S$ to the states.

Let us introduce some notation. For any $s \in S$, $\mathbf{1}_s$ is the probabilistic measure satisfying $\mathbf{1}_s(\{s\}) = 1$, and $\delta(s) = \{\mu : (s, \mu) \in R\}$. A state $s$ is *deterministic* if $\delta(s) = \{\mathbf{1}_{s'}\}$ for some $s' \in S$, *free of unquantified nondeterminism* if $\delta(s) = \{\mu\}$ for some probability measure $\mu$, and *unquantified* if for every $\mu \in \delta(s)$ there is a $s'$ such that $\mu = \mathbf{1}_{s'}$. A PARS is deterministic, free of unquantified nondeterminism, or unquantified if all of its states are so. It is *discrete* if $S$ is finite, $\delta(s)$ is finite for all $s \in S$, and every $\mu \in \delta(s)$ has a finite domain. A PARS that is free of unquantified nondeterminism can be identified with a Markov process, or with a discrete-time Markov chain if it is also discrete. Any discrete PARS can be identified with a Markov decision process provided an appropiate labeling of the options.

**Definition 2** (PARS for $\mathcal{XS}$)**.** Given a term $t$ and a strategy expression $\alpha$, let $\mathcal{XS}_{t,\alpha}$ be $\{q \in \mathcal{XS} : t \,@\, \alpha \to_{s,c}^* q \wedge \mathrm{valid}(q)\}$ endowed of an appropriate $\sigma$-algebra.[1] The PARS on $\mathcal{XS}$ for $t$ and $\alpha$ is $(\mathcal{XS}_{t,\alpha}, R)$ where $R$ is defined as follows. For any state $q$ of the form $t \,@\, \alpha\, s$ where the outermost combinator in $\alpha$ is standard, we consider the transitions $(q, \mathbf{1}_{q'})$ for any small step $q \to_{s,c} q'$ such that $\mathrm{valid}(q)$. Whenever $\alpha$ is a probabilistic combinator, we define a single transition $(q, \mu)$ where $\mu$ is given by

- For `choice(`$w_1$ `:` $\alpha_1$`, ...,` $w_n$ `:` $\alpha_n$`)`, we define $\mu(\{t \,@\, \alpha_k\, s\}) = \hat{w}_k / W$ where $W = \sum_{i=1}^{n} \hat{w}_k$ and $\hat{w}_k = \theta(w_k)$ if valid$(t \,@\, \alpha_k\, s)$ and $\hat{w}_k = 0$ otherwise.

- For `matchrew` combinators, we define a similar measure with domain on the successors $q_\sigma$ of the nondeterministic rule for each match $\sigma$. The probabilities are given by $\mu(\{q_\sigma\}) = \hat{w}_\sigma / W$ where $W = \sum_{\sigma \in \mathsf{mcheck}(P,t,C,\theta)} \hat{w}_\sigma$ and $\hat{w}_\sigma = \sigma(w)$ if valid$(q_\sigma)$ and $\hat{w}_\sigma = 0$ otherwise. In the case of `xmatchrew` and `amatchrew`, successors $q_{(\sigma,c)}$ are indexed by both matching substitution $\sigma$ and context $c$.

- For `sample` $X$ `:=` $\pi$`(`$t_1$`, ...,` $t_n$`)` `in` $\alpha$, if $\nu : \mathbb{R} \to [0,1]$ is the measure of the distribution $\pi(\theta(t_1), \ldots, \theta(t_n))$, we have that

$$\mu(\{t \,@\, \alpha\, \theta[X/x]\, s : x \in U\}) = \nu(U)/W$$

  for any measurable set $U \subseteq \operatorname{dom} \nu \subseteq \mathbb{R}$ where $W = \nu(\{x \in \operatorname{dom} \nu : \text{valid}(t \,@\, \alpha\, \theta[X/x]\, s)\})$.

We know that $W > 0$ in every case by the validity of the states.

Finally, the probabilistic transitions for composed execution states, like rewc and subterm, are derived according to their nondeterministic semantic rules in Figures 2 and 3 of [4]. States that are rewritten with rules without premises are handled as in the $t \,@\, \alpha\, s$ case for $\alpha$ headed by a non-probabilistic combinator. States to be reduced by rules with premises are handled in the straighforward way, i.e. there is a transition $(q, \mu)$ from $q$, there is also a transition $(Q(q), \mu')$ from $Q(q)$ where $\mu'(\{Q(q') : q \in U\}) = \mu(U)$ for any $U \subseteq \operatorname{dom} \mu$.

For example, consider a rewrite system with states $S = \{\mathtt{a}, \mathtt{b}, \mathtt{c}\}$ and two rules `ab` and `ac` that rewrite the first to the second letter. The PARS for the strategy $\alpha = \mathtt{a} \,@\, \mathtt{ab} \,|\, \beta$ with $\beta =$ `choice(2 : ab, 3 : ac, 5 : fail)` is $(\mathcal{XS}_{\mathtt{a},\alpha}, R)$ with

$$R = \{(\mathtt{a} \,@\, \alpha, \{\mathbf{1}_{\mathtt{b} \,@\, \varepsilon}\}), (\mathtt{a} \,@\, \alpha, \{\mathbf{1}_{\mathtt{a} \,@\, \beta}\}), (\mathtt{a} \,@\, \beta, \{\mu\}), (\mathtt{a} \,@\, \mathtt{ab}, \{\mathbf{1}_{\mathtt{b} \,@\, \varepsilon}\}), (\mathtt{a} \,@\, \mathtt{ac}, \{\mathbf{1}_{\mathtt{c} \,@\, \varepsilon}\})\}$$

where $\mu(\{\mathtt{a} \,@\, \mathtt{ab}\}) = 2/5$ and $\mu(\{\mathtt{a} \,@\, \mathtt{ac}\}) = 3/5$. As a consequence, $\delta(\mathtt{a} \,@\, \alpha) = \{\mathbf{1}_{\mathtt{b} \,@\, \varepsilon}, \mathbf{1}_{\mathtt{a} \,@\, \beta}\}$, $\delta(\mathtt{a} \,@\, \beta) = \{\mu\}$, $\delta(\mathtt{a} \,@\, \mathtt{ab}) = \{\mathbf{1}_{\mathtt{b} \,@\, \varepsilon}\}$, $\delta(\mathtt{a} \,@\, \mathtt{ac}) = \{\mathbf{1}_{\mathtt{c} \,@\, \varepsilon}\}$, and $\delta(q) = \emptyset$ otherwise. This PARS is neither deterministic, nor free of unquantified nondeterminism, nor unquantified, but it is discrete.

In the following, we will construct discrete structures, so let us assume that the PARS for the given term $t$ and strategy $\alpha$ is discrete. Our goal is to construct a probabilistic graph whose transitions are one-step rule applications allowed by the strategy, for what we have to collapse control and a system transitions to $\twoheadrightarrow \,=\, \to_s \circ \to_c^*$. For the next definition, we also assume that it is free of unquantified nondeterminism.

**Definition 3** (DTMC induced by a strategy)**.** The discrete-time Markov chain induced by a strategy $\alpha$ that is free of unquantified nondeterminism is given by $(X, P, P_0)$ where $X = X_0 \cup \{t' \,@\, \varepsilon : q \to_c^* t' \,@\, \varepsilon, q \in X_0\}$, $X_0 = \{q : t \,@\, \alpha \twoheadrightarrow^* q\}$, $P_0(t \,@\, \alpha) = 1$, and

$$P(q, q') = \lim_{n \to \infty} \sum_{q_1 \cdots q_m \in \mathrm{Path}^n_{\twoheadrightarrow}(q,q')} \mu_1(\{q_2\}) \cdots \mu_{m-1}(\{q_m\}) \qquad q' \in X_0$$

$$P(q, \mathrm{cterm}(q) \,@\, \varepsilon) = 1 - \sum_{q' \in X_0} P(q, q')$$

where $\mathrm{Path}^n_{\twoheadrightarrow}(q, q') = \{q_1 q_2 \cdots q_m : q = q_1 \to_c q_2 \to_c \cdots \to_c q_{n-1} \to_s q_m = q', m \leq n\}$ and $\mu_k$ is the only element of $\delta(q_k)$ for the PARS in Definition 2.

The formula for $P(q, q')$ sums the probabilities $\mu_1(\{q_2\}) \cdots \mu_{m-1}(\{q_m\})$ of every path from $q$ to $q'$ expanding a $\twoheadrightarrow$ transition. However, since these paths may be infinitely many, the sum is calculated as a limit of length-bounded probabilities. Notice that this DTMC is well-defined since the non-decreasing sequence whose limit is calculated is bounded above by 1. Moreover, the assignment of the remaining probability to the "solution state" $\text{cterm}(q) \to \varepsilon$ is legitimate, since the only option apart from reaching a state $q'$ such that $q \twoheadrightarrow q'$ is $q \to_c^* \text{cterm}(q) \,@\, \varepsilon$ by the definition of valid. In effect, while there may be nonterminating chains of $\to_c$ transitions in some pathological cases, they would have probability zero because they will take infinitely many transitions of probability $p < 1$ since some states must have a different successor of positive probability for reaching a solution or a $\to_s$ transition.

Now, we allow for unquantified nondeterminism in order to derive Markov decision processes. Given any PARS and a choice $\lambda : S^+ \to \mathcal{M}(S)$ of probability measures $\lambda(ws) \in \delta(s)$ for every finite path $ws$ in $S$, we can calculate a single probability $P(w, \lambda)$ for any path $w \in S^*$. We define $P(\varepsilon, \lambda) = 1$, $P(s, \lambda) = 1$, and $P(ws, \lambda) = P(w, \lambda) \cdot \lambda(w)(\{s\})$ for $w \in S^+$ and $s \in S$.

**Definition 4** (MDP induced by a strategy)**.** The Markov decision process induced by the strategy $\alpha$ whose PARS does not contain cycles of $\to_c$ transitions is given by $(X, A, P, P_0)$ where $X$ and $P_0$ are defined as in Definition 3, $A \subseteq \{1, \dots, N\}$ with $N = \Pi_{q' \in X} |\delta(q')| \cdot \max_{q \in X_0} |\text{Paths}_{\to_c^*}(q, q')|$, and $P$ is defined as follows. For each choice $\lambda$, consider

$$P_{\lambda,q}(q') = \lim_{n\to\infty} \sum_{w \in \text{Path}^n_{\twoheadrightarrow}(q,q')} P(w, \lambda) \qquad\qquad q' \in X_0$$

$$P_{\lambda,q}(\text{cterm}(q) \,@\, \varepsilon) = 1 - \sum_{q' \in X_0} P_{\lambda,q}(q')$$

Then, enumerate the elements of the set $\{P_{\lambda,q}\} = \{P_{q,1}, \dots, P_{q,n}\}$ and take $P(q, k, q') = P_{q,k}(q')$.

The construction of the MDP is simpler when the discrete PARS for $\mathcal{XS}$ satisfies the *well-behaved nondeterminism* property: for every path $q_1 \cdots q_n$ from a state $q$ to state $q'$ such that $q \twoheadrightarrow q'$ there is an index $k$ such that $q_1, \dots, q_{k-1}$ are unquantified and $q_{k+1}, \dots, q_n$ are free of unquantified nondeterminism (notice that $q_k$ is arbitrary). In other words, when the first quantified choice appears no other unquantified choice is possible.

# Notes

[1]Assuming that Maude's `Float` type is a 64-bit floating-point number, we can take $\mathbb{R} \equiv \{0,1\}^{64}$ and use the discrete $\sigma$-algebra $\mathcal{P}(\mathcal{XS})$ on $\mathcal{XS}$. Consequently, the continuous probabilistic distributions in the `sample` operator should be discretized as they are in their implementations. Otherwise, we may consider `Float` terms as actual real numbers. Given an execution state it has a finite number of real values in its variable contexts, say $q(v_1, \dots, v_n)$ for some $v_1, \dots, v_n \in \mathbb{R}$. We can take the $\sigma$-algebra consisting of the sets $\{q(x_1, \dots, x_n) : x \in U\}$ for every execution state pattern $q$ and every measurable set $U \subseteq \mathbb{R}^n$.

# References

[1] Oana Andrei, Maribel Fernández, Hélène Kirchner, Guy Melançon, Olivier Namet, and Bruno Pinaud. PORGY: strategy-driven interactive transformation of graphs. In Rachid Echahed, editor, *Proceedings 6th International Workshop on Computing with Terms and Graphs, TERMGRAPH 2011, Saarbrücken, Germany, 2nd April 2011*, volume 48 of *EPTCS*, pages 54–68, 2011. doi: 10.4204/EPTCS.48.7.

[2] Olivier Bournez and Florent Garnier. Proving positive almost sure termination under strategies. In Frank Pfenning, editor, *Term Rewriting and Applications, 17th International Conference, RTA 2006, Seattle, WA, USA, August 12-14, 2006, Proceedings*, volume 4098 of *Lecture Notes in Computer Science*, pages 357–371. Springer, 2006. DOI: 10.1007/11805618_27. URL: https://doi.org/10.1007/11805618%5C_27.

[3] Olivier Bournez and Claude Kirchner. Probabilistic rewrite strategies. Applications to ELAN. In Sophie Tison, editor, *RTA 2002*, volume 2378 of *LNCS*, pages 252–266. Springer, 2002. DOI: 10.1007/3-540-45610-4_18.

[4] Rubén Rubio, Narciso Martí-Oliet, Isabel Pita, and Alberto Verdejo. The semantics of the Maude strategy language. Technical report 01/21, Departamento de Sistemas Informáticos y Computación, Universidad Complutense de Madrid, 2021. URL: https://eprints.ucm.es/67449/.