

Probabilistic Abstract Interpretation: From Trace Semantics to DTMC's and Linear Regression

Alessandra Di Pierro¹ and Herbert Wiklicky²(✉)

¹ Dipartimento di Informatica, Università di Verona, Verona, Italy
alessandra.dipierro@univr.it

² Department of Computing, Imperial College London, London, UK
herbert@doc.ic.ac.uk

Abstract. In order to perform *probabilistic program analysis* we need to consider probabilistic languages or languages with a probabilistic semantics, as well as a corresponding framework for the analysis which is able to accommodate probabilistic properties and properties of probabilistic computations. To this purpose we investigate the relationship between three different types of probabilistic semantics for a core imperative language, namely Kozen's Fixpoint Semantics, our Linear Operator Semantics and probabilistic versions of Maximal Trace Semantics. We also discuss the relationship between Probabilistic Abstract Interpretation (PAI) and statistical or linear regression analysis. While classical Abstract Interpretation, based on Galois connection, allows only for worst-case analyses, the use of the Moore-Penrose pseudo inverse in PAI opens the possibility of exploiting statistical and noisy observations in order to analyse and identify various system properties.

1 Introduction

In this contribution we will address a topic which we believe is dear to the hearts of Hanne and Flemming, namely Abstract Interpretation based techniques in program analysis [1–4]. We will concentrate on the treatment of the probabilistic setting where either the program or its semantics or both contain an element of chance that can be used to refine the possible nondeterminism associated with their models. As program analysis is essentially based on the semantics of programs, we will first describe three different probabilistic semantics that could be used as a basis for probabilistic analysis by clarifying the differences and relationship between them, and discussing their potential for the construction of precise program analyses. As a result of this comparison it will be clear that the Probabilistic Abstract Interpretation framework originally introduced in [5, 6] is not an instance of a probabilistic application of classical Abstract Interpretation as recently suggested in [7] in order to analyse probabilistic programs.

The use of linear operators on vector spaces – more concretely on Hilbert spaces – for the definition of a probabilistic semantics is an important feature

of the Probabilistic Abstract Interpretation framework for several reasons: (i) it provides a well-defined notion of generalised inverse that enjoys properties similar to the concretisation/abstraction functions in the Galois connection framework; (ii) it allows us to exploit a well-defined metric (the Euclidean distance) in order to achieve quantitative results for our static analyses; (iii) it is an appropriate setting where statistical models can be used to enhance the power of static analysis techniques with information gathered via observations.

While we have variously addressed the first two points in our previous work, the potentiality of Probabilistic Abstract Interpretation for performing a kind of *statistical* program analysis was never completely explored before. As another result we will show in this paper that, contrary to the typical computer scientist approach that constructs observations from models, it is sometimes useful to define a model starting from observations, as typically done in statistics. To this end, the particular notion of generalised inverse defining Probabilistic Abstract Interpretation – namely the Moore-Penrose pseudo-inverse [8–10] – makes it very natural to use statistical techniques such as linear regression [10, 11] for constructing abstractions that are as close as possible to the actual system with respect to the observed behaviour.

2 Probabilistic Semantics

There exist a number of proposals for probabilistic languages. These can be based on procedural languages, e.g. [12–14], functional ones, e.g. [15–17], but also declarative ones, like [5, 18]. Besides this there is also a substantial work in probabilistic process algebras [19, 20]. It would be impossible to discuss or even to mention all these approaches here in detail, so we will only concentrate on a small (core) procedural language, which will call **pWhile** and which is essentially the same as the one in [12].

Similarly, a number of approaches have been proposed for defining a semantics for probabilistic programs, not least in order to allow for some form of static program analysis. Usually, it is straightforward to define an operational semantics for a probabilistic extension of a deterministic language; this can be achieved for example by replacing the original (unlabelled) transition relation of an SOS semantics with a weighted version, where the weights represent the probabilities associated with random choices or assignments. Some arguably more useful kinds of semantics are, for example, Kozen’s Fixed-Point Semantics (KFS) [12], the Linear Operator Semantics (LOS) introduced by the authors in [21], and the probabilistic Maximal Trace Semantics (MTS) of [7]. We will concentrate in the following on these three models but again stress that many other approaches exist, which are based e.g. on domain theory [22–24], weakest preconditions [25, 26], and the monadic approach in [16, 27].

2.1 A Probabilistic Language

The syntax of the language we consider is a straightforward extension of an imperative language with a probabilistic assignment “ $x \text{ ?}=\ \rho$ ” where ρ repre-

sents a probability distribution on the set **Value** of possible values of x which associates to every $v_i \in \mathbf{Value}$ a probability p_i . As usual we require for distributions that $0 \leq p_i \leq 1$ and $\sum_i p_i = 1$; these probabilities are all constant, i.e. we do not consider here dynamical changes of distributions. For so-called sub-probability distributions we require $0 \leq p_i \leq 1$ but only $\sum_i p_i \leq 1$. We denote (sub-)probability distributions by sets of pairs $\{\langle v_i, p_i \rangle\}_i$ which express the fact that a constant value v_i has probabilities p_i (pairs with probability $p_i = 0$ can be omitted).

The syntax of statements is given below. We also provide a labelled version of this syntax (cf. [4]) in order to be able to refer to certain program points in a program analysis context. For details on (arithmetic) expressions $f(x_1, \dots, x_n)$ (sometimes denoted simply by e or a) and (Boolean) expressions or tests b , etc. we refer to e.g. [4, 14].

$ \begin{array}{l} S ::= \text{skip} \\ \quad \quad x := f(x_1, \dots, x_n) \\ \quad \quad x \text{ ?}=\rho \\ \quad \quad S_1; S_2 \\ \quad \quad \text{if } b \text{ then } S_1 \text{ else } S_2 \text{ fi} \\ \quad \quad \text{while } b \text{ do } S \text{ od} \end{array} $	$ \begin{array}{l} S ::= [\text{skip}]^\ell \\ \quad \quad [x := f(x_1, \dots, x_n)]^\ell \\ \quad \quad [x \text{ ?}=\rho]^\ell \\ \quad \quad S_1; S_2 \\ \quad \quad \text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2 \text{ fi} \\ \quad \quad \text{while } [b]^\ell \text{ do } S \text{ od} \end{array} $
---	--

It would also be possible to allow for a probabilistic choice construct of the form “choose $p_1 : S_1$ or $p_2 : S_2$ ro”, but in order to keep things simple we omit it in our treatment. This statement can be implemented, for example, as $c \text{ ?}=\rho$; if $c == 0$ then S_1 else S_2 fi with $\rho = \{\langle 0, p_1 \rangle, \langle 1, p_2 \rangle\}$. Further details on the (intuitive and operational) semantics of this language can be found for example in [14, 21, 28].

Though we only deal with constant probabilities in the following we will implicitly always normalise probabilities in a distribution (we cannot assume that a programmer provides the correct probabilities), and we will only allow for rational values (non-rational real values for p_i raise issues of computability we will avoid). This means that we can also require that the p_i are integers indicating the probability ratio between different alternatives.

Example 1. We will consider the following **pWhile** program as a running example throughout the paper (its labelled version can be found below in Example 5):

```

while true do
  if (x == 1)
    then x ?= {⟨0, p⟩, ⟨1, 1 - p⟩}
    else x ?= {⟨0, 1 - q⟩, ⟨1, q⟩}
  fi
od
    
```

This program may be thought of implementing a scheduler in some protocol where $x \mapsto 0$ and $x \mapsto 1$ determines which of two processes has, for example, control over a communication channel.

Clearly the execution of this program never terminates: a random switching between the state $x \mapsto 0$ and $x \mapsto 1$ is performed indefinitely according to the probabilities p and q .

In the following we will assume that the state space (and thus the set of configurations) is finite. This makes the treatment substantially simpler as we can avoid topological and measure theoretic details (for which we refer to [28]) and work with just linear algebraic notions instead of functional analytical or operator algebraic ones, cf. [29, 30] etc. This finiteness condition is fulfilled by the example above. It should be noted that the finiteness of the state space however still allows for infinite executions.

2.2 Kozen's Fixed-Point Semantics (KFS)

A well-known denotational semantics for probabilistic programs was introduced by Kozen in the 1980s [12] based on bounded Banach space operators. This is a fixed-point I/O semantics that describes how an input probability distribution (or in general a measure) is transformed into an output sub-probability distribution/measure. It only records contributions of terminating processes. The probabilities of non-terminating, i.e. infinite, computations “gets lost” so the final outcome is no longer normalised or a full probability distribution/measure. As a consequence the semantics of all non-terminating processes is the same (cf. also [28]).

In Kozen's language in [12] the element of chance is introduced via random assignments. In the semantical interpretation of this language, all the actual executions of a program are however deterministic, as all possible choices are made beforehand [12, Section 3.2.2, p336]. More precisely, before the execution of a program commences, all later probabilistic choices have already been resolved by picking an $\omega \in \Omega$ with $(\Omega, \mathcal{E}, \mu)$ an appropriate measure space (\mathcal{E} the σ -algebra of measurable events and μ a probability measure). The semantics of a program is then parametric in this event or scenario ω which determines the probability that the otherwise deterministic execution of a program may effectively happen.

Example 2. Consider the following simple program:

$$x \text{ ?} = \{ \langle 0, \frac{1}{3} \rangle, \langle 1, \frac{2}{3} \rangle \}; \quad x \text{ ?} = \{ \langle 0, \frac{1}{2} \rangle, \langle 1, \frac{1}{2} \rangle \}.$$

The minimal event space we need for defining a semantics for this program is $\Omega = \{0, 1\} \times \{0, 1\}$ and, because this is a finite set, we can take the whole power-set $\mathcal{E} = \mathcal{P}(\Omega)$ as the σ -algebra of measurable sets. The probability measure of the elements in Ω is then: $\mu(\{(0, 0)\}) = \frac{1}{6}$, $\mu(\{(0, 1)\}) = \frac{1}{6}$, $\mu(\{(1, 0)\}) = \frac{1}{3}$, and $\mu(\{(1, 1)\}) = \frac{1}{3}$.

After a scenario ω has been picked, the program behaves exactly as one of the following deterministic programs:

$$\begin{aligned} \text{for } \omega = (0, 0) \text{ we execute } & \text{''}x := 0; x := 0\text{''} \text{ with probability } \frac{1}{6}, \\ \text{for } \omega = (0, 1) \text{ we execute } & \text{''}x := 0; x := 1\text{''} \text{ with probability } \frac{1}{6}, \\ \text{for } \omega = (1, 0) \text{ we execute } & \text{''}x := 1; x := 0\text{''} \text{ with probability } \frac{1}{3}, \\ \text{for } \omega = (1, 1) \text{ we execute } & \text{''}x := 1; x := 1\text{''} \text{ with probability } \frac{1}{3}. \end{aligned}$$

In the Kozen semantics we can identify a state with a distribution on \mathbf{Value}^n , where n is the number of variables and \mathbf{Value} is the set of possible values of a variable which we assume here – as said before – to be finite. Thus, a probabilistic state (as a distribution $\sigma \in \mathcal{D}(\mathbf{Value}^n)$) can be seen as a normalised element (in the sense of the 1-norm) in the vector space $\mathcal{V}(\mathbf{Value}^n)$. The space $\mathcal{V}(X)$, which allows for the representation of distributions as well as sub-distributions on X , is defined as the set of linear combinations of elements in X , i.e.

$$\mathcal{V}(X) = \left\{ \sum_i \lambda_i x_i \mid x_i \in X \wedge \lambda_i \in \mathbb{R} \right\}.$$

This space is isomorphic to $\mathbb{R}^{|X|}$ with $|X|$ the cardinality of X . Vector addition and scalar product are defined pointwise. We can identify $x_i \in X$ with the base vectors of $\mathcal{V}(X)$ and any element in $\mathcal{V}(X)$ with its coordinates, i.e. the tuple $(\lambda_i)_i$. This space is equipped with an inner product $\langle (\lambda_i)_i | (\nu_i)_i \rangle = \sum_i \lambda_i \nu_i$ and one of various norms, e.g. $\|(\lambda_i)_i\|_1 = \sum_i |\lambda_i|$ and $\|(\lambda_i)_i\|_2 = \sqrt{\sum_i |\lambda_i|^2} = \sqrt{\langle (\lambda_i)_i | (\lambda_i)_i \rangle}$. The choice of one norm or another is nevertheless largely irrelevant in the finite dimensional case where all norms are equivalent. In fact, the topology on finite dimensional vector spaces is uniquely determined by the algebraic structure, cf. e.g. [31, 1.22].

The Kozen semantics of a program P is then given by the linear operator $\llbracket P \rrbracket_{KFS} \in \mathcal{L}(\mathcal{V}(\mathbf{Value}^n))$ where $\mathcal{L}(X)$ is the set of linear maps \mathbf{T} on X , i.e. $\mathbf{T}(x + y) = \mathbf{T}(x) + \mathbf{T}(y)$ and $\mathbf{T}(\lambda x) = \lambda \mathbf{T}(x)$:

$$\llbracket P \rrbracket_{KFS} : \mathcal{V}(\mathbf{Value}^n) \rightarrow \mathcal{V}(\mathbf{Value}^n),$$

which is the solution to the following set of equations:

$$\begin{aligned} \llbracket \text{skip} \rrbracket_{KFS} &= \mathbf{I} \\ \llbracket x := f(x_1, \dots, x_n) \rrbracket_{KFS} &= \mathbf{U}(x \leftarrow f(x_1, \dots, x_n)) \\ \llbracket x \text{ ?} = \rho \rrbracket_{KFS} &= \sum_v \rho(v) \mathbf{U}(x \leftarrow v) \\ \llbracket S_1; S_2 \rrbracket_{KFS} &= (\llbracket S_1 \rrbracket_{KFS} \llbracket S_2 \rrbracket_{KFS}) \\ \llbracket \text{if } b \text{ then } S_1 \text{ else } S_2 \text{ fi} \rrbracket_{KFS} &= (\mathbf{P}(b) \llbracket S_1 \rrbracket_{KFS} + \mathbf{P}(-b) \llbracket S_2 \rrbracket_{KFS}) \\ \llbracket \text{while } b \text{ do } S \text{ od} \rrbracket_{KFS} &= (\mathbf{P}(b) \llbracket S \rrbracket_{KFS} \llbracket \text{while } b \text{ do } S \text{ od} \rrbracket_{KFS} + \mathbf{P}(-b)). \end{aligned}$$

The operator \mathbf{I} is the identity on $\mathcal{V}(\mathbf{Value}^n)$ represented by a matrix with $(\mathbf{I})_{vv} = 1$ and 0 otherwise for $v = (v_1, \dots, v_n) \in \mathbf{Value}^n$. The matrix representation of the test or projection operators \mathbf{P} is given by a diagonal matrix with $(\mathbf{P}(b))_{vv} = 1$ if $b(v)$ holds for $v \in \mathbf{Value}^n$ and 0 otherwise. Note that $\mathbf{P}(\text{true}) = \mathbf{I}$ and that $\mathbf{P}(-b) = \mathbf{I} - \mathbf{P}(b)$. The assignment or update operator \mathbf{U} is given by a matrix with entries $(\mathbf{U}(x_i \leftarrow f(x_1, \dots, x_n)))_{v, F(v)} = 1$ for all $v \in \mathbf{Value}^n$ and 0 otherwise, where $F : \mathbf{Value}^n \rightarrow \mathbf{Value}^n$ is defined as

$$F(v_1, \dots, v_{i-1}, v_i, v_{i+1}, \dots, v_n) = (v_1, \dots, v_{i-1}, f(v_1, \dots, v_n), v_{i+1}, \dots, v_n).$$

This definition is equivalent to that given in [12, p339]).

The *existence* of a solution to these equations is guaranteed in general (i.e. also for infinite state spaces) by the Brouwer-Schauders fixed-point theorem (see e.g. [30, 32]). The least fixed-point can be *constructed* iteratively via a “super-operator” $\tau : \mathcal{L}(\mathcal{V}(\mathbf{Value}^n)) \rightarrow \mathcal{L}(\mathcal{V}(\mathbf{Value}^n))$ which encodes the above equations and by exploiting the lifted point-wise order on distributions/measures.

Example 3. Consider again the program P in Example 1. As no executions of this program will ever terminate, there is no proper (sub-)probability distribution describing the final state. Thus Kozen’s semantics, which describes the I/O behaviour, is trivial:

$$\llbracket P \rrbracket_{KFS} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix} = \mathbf{O}$$

i.e. the zero operator $\llbracket P \rrbracket_{KFS} : \mathcal{V}(\{x \mapsto 0, x \mapsto 1\}) \rightarrow \mathcal{V}(\{x \mapsto 0, x \mapsto 1\})$.

This is also justified by the fixed-point construction described in [12, p 341]. The semantics of the statement S given by

if $(x == 1)$ then $x \text{ ?= } \{\langle 0, p \rangle, \langle 1, 1 - p \rangle\}$ else $x \text{ ?= } \{\langle 0, 1 - q \rangle, \langle 1, q \rangle\}$ fi

forming the body of the loop is easily computed as:

$$\llbracket S \rrbracket_{KFS} = \begin{pmatrix} p & 1 - p \\ 1 - q & q \end{pmatrix},$$

but whatever the semantics $\llbracket S \rrbracket_{KFS}$ of the body of loop is, the Kozen semantics of the whole program P is the (appropriate) supremum of a sequence of matrices $\tau^n(\mathbf{O})$ with $n = 1, 2, 3, \dots$ (starting with the zero matrix \mathbf{O}):

$$\tau^n(\mathbf{O}) = \sum_{k=0}^{n-1} (\mathbf{P}(\mathbf{true}) \llbracket S \rrbracket_{KFS})^k \mathbf{P}(\mathbf{false}) = \sum_{k=0}^{n-1} (\mathbf{I} \llbracket S \rrbracket_{KFS})^k \mathbf{O} = \mathbf{O}.$$

That is, for all $n = 1, 2, 3, \dots$ we have $\tau^n(\mathbf{O}) = \mathbf{O}$ and thus $\llbracket P \rrbracket_{KFS} = \mathbf{O}$.

We also represent (sub-)probability distributions as row vectors; the application of an operator or linear map $\mathbf{T}(x)$ is thus expressed by *post-multiplication* $x \cdot \mathbf{T}$ rather than *pre-multiplication* as it can be found elsewhere (e.g. [12]).

Example 3 describes the situation of a program that never terminates on all inputs. More interestingly, Kozen’s semantics also allows us to model programs that terminate with probability $0 < p < 1$, as shown in the following example.

Example 4. Consider the programs Q , Q' and Q'' which incorporate the program P in Example 3:

if $(x == 1)$	$x := 1;$	$x \text{ ?= } \{\langle 0, \frac{1}{2} \rangle, \langle 1, \frac{1}{2} \rangle\};$
then $x := 0$	if $(x == 1)$	if $(x == 1)$
else P	then $x := 0$	then $x := 0$
fi	else P	else P
	fi	fi

The operator $\llbracket Q \rrbracket_{KFS}$ of the first program can be easily computed (based on $\llbracket P \rrbracket_{KFS}$ in Example 3). We have the Kozen semantics of the two branches of the `if` statements:

$$\llbracket x := 0 \rrbracket_{KFS} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix} \quad \llbracket P \rrbracket_{KFS} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}$$

as well as for the tests guarding the `if` statement:

$$\llbracket x = 0 \rrbracket_{KFS} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \quad \llbracket x = 1 \rrbracket_{KFS} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} = \mathbf{I} - \llbracket x = 0 \rrbracket_{KFS}$$

Thus by the fifth equation in the definition of the KFS (or section (3.3.4) in [12, p 340]) we get:

$$\llbracket Q \rrbracket_{KFS} = \llbracket x = 1 \rrbracket_{KFS} \llbracket x := 0 \rrbracket_{KFS} + \llbracket x = 0 \rrbracket_{KFS} \llbracket P \rrbracket_{KFS} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}$$

This means that if we have an initial state $\sigma = (p, 1 - p)^t$ which describes the fact that the initial value of x is zero with probability p , and one with probability $1 - p$, then $\sigma \llbracket Q \rrbracket_{KFS} = (p, 0)^t$ (where $.^t$ denotes vector transposition). This is in general (unless $p = 1$) only a sub-probability distribution expressing the fact that this program will terminate with probability p with a zero value for x and that with probability $1 - p$ we have non-termination.

If we consider instead the second program Q' then we have

$$\llbracket x := 1 \rrbracket_{KFS} = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and thus} \quad \llbracket Q' \rrbracket_{KFS} = \llbracket x := 1 \rrbracket_{KFS} \llbracket Q' \rrbracket_{KFS} = \begin{pmatrix} 1 & 0 \\ 1 & 0 \end{pmatrix}$$

That means that independently of the initial value of x we always get (i.e. with probability one) a termination and a zero value for x .

Finally, if we consider the program Q'' we get

$$\llbracket Q'' \rrbracket_{KFS} = \left(\frac{1}{2} \llbracket x := 0 \rrbracket_{KFS} + \frac{1}{2} \llbracket x := 1 \rrbracket_{KFS} \right) \llbracket Q \rrbracket_{KFS} = \begin{pmatrix} \frac{1}{2} & 0 \\ \frac{1}{2} & 0 \end{pmatrix}$$

Here we terminate (again with the resulting x being zero) with a half probability, independently from the initial value of x .

2.3 Linear Operator Semantics (LOS)

The Linear Operator Semantics in [21, 28] constructs the generator of a Discrete Time Markov Chain (DTMC) in a syntax directed fashion. Like Kozen's semantics we can represent the LOS as an operator on the vector space of probabilistic states. However, differently from Kozen's semantics, the definition of this operator is based on the syntax rather than on a denotational domain. Moreover, in order to provide a suitable base for static analysis, we do not construct the LOS of a program by simply translating the SOS transition relation into a DTMC

generator. Instead, we define it in a *structured* way by composing the operators associated with each syntactic elementary components of the program by means of the tensor (or Kronecker) product operation “ \otimes ” on vector spaces (cf. e.g. [33, 34] or [21]).

The state space is constructed starting from the classical states, i.e. states that associate concrete values in $v_i \in \mathbf{Value}$ to variables $x_i \in \mathbf{Var} = \{x_1, \dots, x_n\}$. The classical state space can therefore be defined as $\mathbf{State} = \mathbf{Var} \rightarrow \mathbf{Value}$ or equivalently $\mathbf{State} = \mathbf{Value}_1 \times \dots \times \mathbf{Value}_n = \mathbf{Value}^n$.

In order to describe the probabilistic state of a computation we consider (probability) distributions over (classical) states again – as in Kozen’s construction – as elements in $\mathcal{V}(\mathbf{Value}^n)$. However, we can use the tensor product operation “ \otimes ” to decompose this probabilistic state space, i.e. $\mathcal{V}(X \times Y) = \mathcal{V}(X) \otimes \mathcal{V}(Y)$ and represent probabilistic states thus as elements in $\mathcal{V}(\mathbf{Value}^n) = \mathcal{V}(\mathbf{Value}_1) \otimes \mathcal{V}(\mathbf{Value}_2) \otimes \dots \otimes \mathcal{V}(\mathbf{Value}_n) = \mathcal{V}(\mathbf{Value})^{\otimes n}$.

The LOS is based on the labelled version of the syntax of **pWhile**. This allows us to record not only the values of all variables but also the current point in the program we are executing, i.e. the “program counter”. Thus, the state space of the corresponding DTMC is a space of configurations which also contain information about the current label. This is defined as the space $\mathbf{Conf} = \mathbf{State} \times \mathbf{Label}$ of distributions in $\mathcal{D}(\mathbf{Conf}) \subseteq \mathcal{V}(\mathbf{Conf}) = \mathcal{V}(\mathbf{State}) \otimes \mathcal{V}(\mathbf{Label}) = \mathcal{V}(\mathbf{Value})^{\otimes n} \otimes \mathcal{V}(\mathbf{Label})$.

The LOS $\llbracket P \rrbracket_{LOS}$ of a program P is then an operator in $\mathcal{L}(\mathcal{V}(\mathbf{Conf}))$ or, more precisely

$$\llbracket P \rrbracket_{LOS} : \mathcal{V}(\mathbf{Value}^n) \otimes \mathcal{V}(\mathbf{Label}) \rightarrow \mathcal{V}(\mathbf{Value}^n) \otimes \mathcal{V}(\mathbf{Label}).$$

It is constructed by means of a set $\{\llbracket P \rrbracket_{LOS}\}$ of linear operators describing local changes (at individual labels) as follows:

$$\llbracket P \rrbracket_{LOS} = \sum \{\llbracket P \rrbracket_{LOS}\} = \sum \{\mathbf{G} \mid \mathbf{G} \in \{\llbracket P \rrbracket_{LOS}\}\}.$$

The $\{\llbracket S \rrbracket_{LOS}\}$ associated to a statement S is given by a set of global and local operators, i.e. $\{\llbracket \cdot \rrbracket_{LOS} : \mathbf{Stmt} \rightarrow \mathcal{P}(\Gamma \cup \Lambda)$. Global operators are linear operators on $\mathcal{V}(\mathbf{Conf})$ i.e. $\Gamma = \mathcal{L}(\mathcal{V}(\mathbf{Value}^n) \otimes \mathcal{V}(\mathbf{Label})) = \mathcal{L}(\mathcal{V}(\mathbf{Conf}))$, and local operators are pairs of operators on $\mathcal{V}(\mathbf{State})$ and labels $\ell \in \mathbf{Label}$, i.e. $\Lambda = \mathcal{L}(\mathcal{V}(\mathbf{Value}^n)) \times \mathbf{Label}$.

Global operators provide information about how the computational state changes at a label as well as the control flow; in other words, they define the label of the next statement to be executed. Local operators represent statements for which a “continuation” is not yet known. In order to transform local operators into global ones, we define a continuation operation $\langle \mathbf{F}, \ell \rangle \triangleright \ell' = \mathbf{F} \otimes \mathbf{E}(\ell, \ell')$ which we extend in the obvious way to sets of operators by $\{\langle \mathbf{F}_i, \ell_i \rangle\} \triangleright \ell' = \{\mathbf{F}_i \otimes \mathbf{E}(\ell_i, \ell')\}$ (for global operators, clearly, we have $\mathbf{G} \triangleright \ell' = \mathbf{G}$). Here, $\mathbf{E}(i, j)$ denotes the matrix unit with $(\mathbf{E}(i, j))_{ij} = 1$ and 0 otherwise.

The set $\{\llbracket S \rrbracket_{LOS}\}$ of operators for a statement S is defined inductively on the syntactic structure of S as follows:

$$\begin{aligned}
 \llbracket [\text{skip}]^\ell \rrbracket_{LOS} &= \{\langle \mathbf{I}, \ell \rangle\} \\
 \llbracket [x := e]^\ell \rrbracket_{LOS} &= \{\langle \mathbf{U}(x \leftarrow e), \ell \rangle\} \\
 \llbracket [x \text{ ?}=\rho]^\ell \rrbracket_{LOS} &= \{\langle \sum_{\langle v, p \rangle \in \rho} p \cdot \mathbf{U}(x \leftarrow v), \ell \rangle\} \\
 \llbracket [S_1; S_2] \rrbracket_{LOS} &= (\llbracket [S_1] \rrbracket_{LOS} \triangleright \text{init}(S_2)) \cup \llbracket [S_2] \rrbracket_{LOS} \\
 \llbracket [\text{if } [b]^\ell \text{ then } S_1 \text{ else } S_2 \text{ fi}] \rrbracket_{LOS} &= \{\langle \mathbf{P}(b), \ell \rangle\} \triangleright \text{init}(S_1) \cup \llbracket [S_1] \rrbracket_{LOS} \cup \\
 &\quad \{\langle \mathbf{P}(b)^\perp, \ell \rangle\} \triangleright \text{init}(S_2) \cup \llbracket [S_2] \rrbracket_{LOS} \\
 \llbracket [\text{while } [b]^\ell \text{ do } S \text{ od}] \rrbracket_{LOS} &= \{\langle \mathbf{P}(b), \ell \rangle\} \triangleright \text{init}(S) \cup \llbracket [S] \rrbracket_{LOS} \cup \{\langle \mathbf{P}(b)^\perp, \ell \rangle\}
 \end{aligned}$$

We use elementary update and test operators \mathbf{U} and \mathbf{P} (and its complement $\mathbf{P}^\perp = \mathbf{I} - \mathbf{P}$) as in Kozen's semantics. However, the tensor product structure allows us to define these operators in a different (although equivalent) way.

For a *single* variable the assignment to a constant value $v \in \mathbf{Value}$ is represented by the operator on $\mathcal{V}(\mathbf{Value})$ given by $\mathbf{U}(v) = 1$ if $v = i$ and 0 otherwise. Testing if a *single* variable satisfies a Boolean test b is achieved by a projection operator on $\mathcal{V}(\mathbf{Value})$ with $(\mathbf{P}(b))_{ii} = 1$ if $b(i)$ holds and 0 otherwise.

We extend these to the multivariable case, i.e. for $|\mathbf{Var}| = n > 1$ by defining the following operators on $\mathcal{V}(\mathbf{Value})^{\otimes n}$:

$$\mathbf{P}(s) = \bigotimes_{i=1}^n \mathbf{P}(x_i = s(x_i)) \qquad \mathbf{P}(e = v) = \sum_{\mathcal{E}(e)=v} \mathbf{P}(s),$$

where $\mathbf{P}(s)$ is for testing if we are in a classical state $s \in \mathbf{Value}^n$ while $\mathbf{P}(e = v)$ checks if an expression e evaluates to a constant v (assuming an appropriate evaluation function $\mathcal{E} : \mathbf{Expr} \rightarrow \mathbf{State} \rightarrow \mathbf{Value}$).

Operators for updating a variable x_k in the context of other variables to a constant v or to the value of an expression e are defined on $\mathcal{V}(\mathbf{Value})^{\otimes n}$ by:

$$\mathbf{U}(x_k \leftarrow v) = \bigotimes_{i=1}^{k-1} \mathbf{I} \otimes \mathbf{U}(v) \otimes \bigotimes_{i=k+1}^n \mathbf{I} \qquad \mathbf{U}(x_k \leftarrow e) = \sum_v \mathbf{P}(e = v) \mathbf{U}(x_k \leftarrow v)$$

As we model the semantics of a program as a DTMC, we need to add a final loop ℓ^* (for ℓ^* a fresh label not appearing already in P) when we consider a complete program. This is because a DTMC never terminates and thus we have to simulate termination by an infinite repetition of the final state. We will therefore use $(\llbracket [P] \rrbracket_{LOS} \triangleright \ell^*) \cup \{\langle \mathbf{I} \otimes \mathbf{E}(\ell^*, \ell^*) \rangle\}$ for the construction of $\llbracket [P] \rrbracket_{LOS}$. In this way we also resolve all open or dangling control flow steps, i.e. we deal ultimately with a set containing only global operators.

Example 5. Consider the labelled version of the program in Example 1

```

while [true]1 do
  if [(x == 1)]2
    then [x ?= {⟨0, p⟩, ⟨1, 1 - p⟩}]3
    else [x ?= {⟨0, 1 - q⟩, ⟨1, q⟩}]4
  fi
od

```

In order to define the LOS of this program we construct the state space as $\mathcal{V}(\{x \mapsto 0, x \mapsto 1\}) = \mathbb{R}^2$ (since we have only one variable we do not need the tensor product for this). The space of configurations is $\mathcal{V}(\{x \mapsto 0, x \mapsto 1\}) \otimes \mathcal{V}(\{1, 2, 3, 4, 5\})$, where label 5 is the label of the additional final loop. We will omit the final label (which in this program we actually never reach) in order to deal with smaller matrices. The set $\{\!\{P\}\!\}_{LOS}$ of P will contain the following operators:

$$\begin{aligned} \{\!\{P\}\!\}_{LOS} = & \{\mathbf{P}(\mathbf{true}) \otimes \mathbf{E}(1, 2), \mathbf{P}(x = 1) \otimes \mathbf{E}(2, 3), \mathbf{P}(x = 1)^\perp \otimes \mathbf{E}(2, 4), \\ & (p \cdot \mathbf{U}(x \leftarrow 0) + (1 - p) \cdot \mathbf{U}(x \leftarrow 1)) \otimes \mathbf{E}(3, 1), \\ & ((1 - q) \cdot \mathbf{U}(x \leftarrow 0) + q \cdot \mathbf{U}(x \leftarrow 1)) \otimes \mathbf{E}(4, 1)\} \end{aligned}$$

The concrete matrices representing the operators on $\mathcal{V}(\{x \mapsto 0, x \mapsto 1\}) \otimes \mathcal{V}(\{1, 2, 3, 4\})$ are of the form

$$\begin{aligned} \{\!\{P\}\!\}_{LOS} = & \left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \mathbf{E}(1, 2), \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \otimes \mathbf{E}(2, 3), \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \otimes \mathbf{E}(2, 4), \right. \\ & \left. \left(\begin{pmatrix} p & 0 \\ p & 0 \end{pmatrix} + \begin{pmatrix} 0 & (1-p) \\ 0 & (1-p) \end{pmatrix} \right) \otimes \mathbf{E}(3, 1), \left(\begin{pmatrix} (1-q) & 0 \\ (1-q) & 0 \end{pmatrix} + \begin{pmatrix} 0 & q \\ 0 & q \end{pmatrix} \right) \otimes \mathbf{E}(4, 1) \right\} \end{aligned}$$

or, explicitly

$$\begin{aligned} \{\!\{P\}\!\}_{LOS} = & \left\{ \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix} \otimes \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix} \otimes \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \right. \\ & \left. \begin{pmatrix} p & (1-p) \\ p & (1-p) \end{pmatrix} \otimes \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \begin{pmatrix} (1-q) & q \\ (1-q) & q \end{pmatrix} \otimes \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} \right\} \end{aligned}$$

The sum of these 8×8 matrices gives the operator $\llbracket P \rrbracket_{LOS}$, i.e. the generator of the corresponding DTMC. By including also the final label $\ell^* = 5$, we obtain a 10×10 matrix, which we depict in the following for the case $p = q = \frac{1}{2}$:

$$\llbracket P \rrbracket_{LOS} = \begin{pmatrix} 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ \frac{1}{2} & 0 & 0 & 0 & 0 & \frac{1}{2} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{pmatrix} \begin{array}{l} \dots x \mapsto 0, \ell = 1 \\ \dots x \mapsto 0, \ell = 2 \\ \dots x \mapsto 0, \ell = 3 \\ \dots x \mapsto 0, \ell = 4 \\ \dots x \mapsto 0, \ell = 5 \\ \dots x \mapsto 1, \ell = 1 \\ \dots x \mapsto 1, \ell = 2 \\ \dots x \mapsto 1, \ell = 3 \\ \dots x \mapsto 1, \ell = 4 \\ \dots x \mapsto 1, \ell = 5 \end{array}$$

The entries of this matrix represent the probability of the configuration (i.e. value of x and current label ℓ) each row and column corresponds to. It is perhaps worth noting that this – as one would expect for a DTMC – is indeed a stochastic matrix (i.e. all row sums are one) representing the SOS transition relation.

There is a close relationship between the KFS and the LOS. For basic blocks B – i.e. (random) assignments, tests and skips – the LOS operator is the same as the KFS operator except for an additional control flow step. That means that $\{\dots [B]^i \dots\}_{LOS} = \{\dots, \langle \llbracket B \rrbracket_{KFS}, i \rangle, \dots\}$ or $\{\dots [B]^i \dots\}_{LOS} = \{\dots, \llbracket B \rrbracket_{KFS} \otimes \mathbf{E}(i, j), \dots\}$ for some label j .

Example 6. For the programs in Example 4 with the following labelling

$$\begin{array}{lll} \text{if } [(x == 1)]^1 & [x := 1]^0; & [x ?= \{ \langle 0, \frac{1}{2} \rangle, \langle 1, \frac{1}{2} \rangle \}]^0; \\ \text{then } [x := 0]^2 & \text{if } [(x == 1)]^1 & \text{if } [x == 1]^1 \\ \text{then } P & \text{then } [x := 0]^2 & \text{then } [x := 0]^2 \\ \text{fi} & \text{then } P & \text{then } P \\ & \text{fi} & \text{fi} \end{array}$$

(the labels of P are as in the previous example shifted by an offset of 2), we can describe the LOS using the KFS operators as follows:

$$\begin{aligned} \llbracket Q \rrbracket_{LOS} &= \{ \llbracket [x = 1] \rrbracket_{KFS} \otimes \mathbf{E}(1, 2), \llbracket [x = 0] \rrbracket_{KFS} \otimes \mathbf{E}(1, 3), \\ &\quad \langle \llbracket [x := 0] \rrbracket_{KFS}, 2 \rangle, \langle \llbracket \text{false} \rrbracket_{KFS}, 3 \rangle \} \cup \llbracket P \rrbracket_{LOS} \\ \llbracket Q' \rrbracket_{LOS} &= \{ \llbracket [x := 1] \rrbracket_{KFS} \otimes \mathbf{E}(0, 1) \} \cup \llbracket Q \rrbracket_{LOS} \\ \llbracket Q'' \rrbracket_{LOS} &= \{ (\frac{1}{2} \llbracket [x := 1] \rrbracket_{KFS} + \frac{1}{2} \llbracket [x := 1] \rrbracket_{KFS}) \otimes \mathbf{E}(0, 1), \} \cup \llbracket Q \rrbracket_{LOS} \end{aligned}$$

where we can re-use the LOS semantics of program P (with shifted labelling):

$$\begin{aligned} \llbracket P \rrbracket_{LOS} &= \{ \llbracket \text{true} \rrbracket_{KFS} \otimes \mathbf{E}(3, 4), \llbracket [x = 1] \rrbracket_{KFS} \otimes \mathbf{E}(4, 5), \llbracket [x = 0] \rrbracket_{KFS} \otimes \mathbf{E}(4, 6), \\ &\quad (p \llbracket [x := 0] \rrbracket_{KFS} + (p - 1) \llbracket [x := 1] \rrbracket_{KFS}) \otimes \mathbf{E}(5, 3), \\ &\quad ((q - 1) \llbracket [x := 0] \rrbracket_{KFS} + q \llbracket [x := 1] \rrbracket_{KFS}) \otimes \mathbf{E}(6, 3) \} \end{aligned}$$

Note that the LOS of the three small programs contain not just global but also local operators, namely $\langle \llbracket x := 0 \rrbracket_{KFS}, 2 \rangle$ and $\langle \llbracket \mathbf{false} \rrbracket_{KFS}, 3 \rangle$. This is because we still have to add a terminal label $\ell^* = 7$ for the construction of the complete DTMC generators $\llbracket Q \rrbracket_{LOS}$, $\llbracket Q' \rrbracket_{LOS}$ and $\llbracket Q'' \rrbracket_{LOS}$. The terminal label can be reached from both branches of the **if** statement labelled 2 and 3. However, as $\llbracket \mathbf{false} \rrbracket_{KFS}$ is **O** this operator (which would correspond to a terminating program P) does not actually contribute to the DTMC generator.

2.4 Maximal Trace Semantics (MTS)

Maximal Trace Semantics for non-probabilistic programs has been discussed in [35, 36] and shown to be the most concrete semantics in a hierarchy of various semantics for (non-)deterministic programs. In [7] the MTS is extended to the probabilistic case.

Similar to the Kozen semantics, the conceptual idea is to ban any probabilistic steps from the actual execution of the program and resolve all probabilistic choices (coin flips, rolling of dices) beforehand. The actual execution of a program is therefore purely (non-)deterministic but parameterised by the results of the “pre-run” choices (cf. [7, p 171]). These outcomes represent the events or scenarios of a probability space Ω , which the execution traces depend on.

Given a set of states Σ , a *trace* $\sigma = s_1 s_2 \dots$ is a finite or infinite sequence of elements $s_i \in \Sigma$. Concatenation of traces is juxtaposition, i.e. for $s \in \Sigma$ we have $s\sigma = s s_1 s_2 \dots$ and for $\sigma_1 = s_{11} \dots s_{1n}$ and $\sigma_2 = s_{21} \dots$ we have $\sigma_1 \sigma_2 = s_{11} \dots s_{1n} s_{21} \dots$. We denote by Σ^+ the set of finite traces, by Σ^* the set $\Sigma^+ \cup \{\varepsilon\}$, where ε is the empty trace of length 0, by Σ^∞ the infinite traces, by $\Sigma^{+\infty}$ the set $\Sigma^+ \cup \Sigma^\infty$ and by $\Sigma^{*\infty}$ the set $\Sigma^* \cup \Sigma^\infty$. For sets of traces X, Y, \dots in $\Sigma^{*\infty}$, we can define the following operations: $X^\infty = X \cap \Sigma^\infty$, $X^+ = X \cap \Sigma^+$, $X|_Y = \{s\sigma_X \in X \mid \exists \sigma : \sigma s \in Y^+\}$, and $X; Y = X^\infty \cup \{\sigma_X s \sigma_Y \mid \sigma_X s \in X^+ \wedge s \sigma_Y \in Y^+\}$.

The MTS is defined as a function $\llbracket S \rrbracket_{MTS} : \mathbf{Stmt} \rightarrow \Omega \rightarrow \mathcal{P}(\Sigma^{+\infty})$ where $\Sigma = \mathbf{State}$. In order to combine non-determinism with probabilities each scenario $\omega \in \Omega$ is associated to a whole set of possible traces. Thus $\llbracket S \rrbracket_{MTS}$ is defined by

$$\begin{aligned} \llbracket \mathbf{skip} \rrbracket_{MTS}(\omega) &= \{ss \mid s \in \Sigma\} \\ \llbracket x := e \rrbracket_{MTS}(\omega) &= \{ss[x \mapsto \llbracket e \rrbracket(\omega)s] \mid s \in \Sigma\} \\ \llbracket S_1; S_2 \rrbracket_{MTS}(\omega) &= \llbracket S_1 \rrbracket_{MTS}(\omega); \llbracket S_2 \rrbracket_{MTS}(\omega) \\ \llbracket b \rrbracket_{MTS}(\omega) &= \{s \mid \llbracket b \rrbracket(\omega)s\} \\ \llbracket \mathbf{if } b \mathbf{ then } S_1 \mathbf{ else } S_2 \mathbf{ fi} \rrbracket_{MTS}(\omega) &= \llbracket b \rrbracket_{MTS}(\omega); \llbracket S_1 \rrbracket_{MTS}(\omega) \cup \llbracket \neg b \rrbracket_{MTS}(\omega); \llbracket S_2 \rrbracket_{MTS}(\omega) \\ \llbracket \mathbf{while } b \mathbf{ do } S \mathbf{ od} \rrbracket_{MTS}(\omega) &= \text{lfp} \lambda X. \llbracket b \rrbracket_{MTS}(\omega) \cup \llbracket \neg b \rrbracket_{MTS}(\omega); \llbracket S \rrbracket_{MTS}(\omega); X \end{aligned}$$

According to the definition in [7, Example 4] the evaluation $\llbracket e \rrbracket$ of an expression e depends on the scenario ω , i.e. $\llbracket e \rrbracket : \Omega \rightarrow (\Sigma \rightarrow \Sigma)$. The language considered in [7] does actually not have either random assignments or a choice construct; the former is instead implemented via a kind of “system call”, i.e. $\mathbf{x} := \mathbf{random}(\rho)$.

We can reformulate the MTS in the case of the **pWhile** language where no non-determinism is present: Once a scenario ω is fixed there is only one trace for every initial state or configuration which is actually executed. We are not interested in the scenarios $\omega \in \Omega$ themselves but only in their probabilities $\mu(\omega)$,

i.e. the probability that a certain trace gets executed. Thus, for a fixed initial state or configuration s the MTS of a program in **pWhile** can be seen as a distribution over traces. The probability for each trace σ is inherited from the scenario ω it depends on. We will use in the following the notation $\{\langle\sigma, \mu(\omega)\rangle\}$ to express that a trace σ is executed with probability $\mu(\omega)$.

We can define the map $\llbracket \cdot \rrbracket_{MTS} : \mathbf{Stmt} \rightarrow \mathcal{V}(\Sigma^{+\infty})$ implicitly, i.e. as solution to the following equations:

$$\begin{aligned}
 \llbracket \text{skip} \rrbracket_{MTS} &= \{\langle ss, 1 \rangle \mid s \in \Sigma\} \\
 \llbracket x := e \rrbracket_{MTS} &= \{\langle ss[x \mapsto \llbracket e \rrbracket s], 1 \rangle \mid s \in \Sigma\} \\
 \llbracket x \text{ ?= } \rho \rrbracket_{MTS} &= \{\langle ss[x \mapsto v], \rho(v) \rangle \mid s \in \Sigma \wedge \rho(v) \neq 0\} \\
 \llbracket S_1; S_2 \rrbracket_{MTS} &= \llbracket S_1 \rrbracket_{MTS}; \llbracket S_2 \rrbracket_{MTS} \\
 \llbracket \text{if } b \text{ then } S_1 \text{ else } S_2 \text{ fi} \rrbracket_{MTS} &= \{\langle s, 1 \rangle \mid \text{for } \llbracket b \rrbracket(s) = \text{true}; \llbracket S_1 \rrbracket_{MTS} \\
 &\quad \cup \{\langle s, 1 \rangle \mid \text{for } \llbracket b \rrbracket(s) = \text{false}; \llbracket S_2 \rrbracket_{MTS}\} \\
 \llbracket \text{while } b \text{ do } S \text{ od} \rrbracket_{MTS} &= \{\langle s, 1 \rangle \mid \text{for } \llbracket b \rrbracket(s) = \text{true}; \llbracket S; \text{while } b \text{ do } S \text{ od} \rrbracket_{MTS} \\
 &\quad \cup \{\langle s, 1 \rangle \mid \text{for } \llbracket b \rrbracket(s) = \text{false};\}
 \end{aligned}$$

Clearly, the evaluation of deterministic functions or expressions is independent of the scenario ω . For random assignments we produce a set of weighted traces, one trace for each $v \in \mathbf{Value}$ with non-vanishing probability according to the distribution ρ . We extend the concatenation operation for traces to probabilistic ones in the obvious way: $\langle X, p_X \rangle; \langle Y, p_Y \rangle = \langle X; Y, p_X p_Y \rangle$ in order to define the semantics of sequential statements. The operation “;” also extends pointwise to sets of weighted traces in $\mathcal{V}(\Sigma^{+\infty})$. The union construction \cup of sets of weighted tuples corresponds to a sum if we take them as elements in the vector space $\mathcal{V}(\Sigma^{+\infty})$.

It should be noted that this formulation of the MTS for a purely probabilistic language eliminates the dependency on the scenarios $\omega \in \Omega$ but not on the initial state $s \in \Sigma$. This means that for a statement S the weighted set of traces $\llbracket S \rrbracket_{MTS} \in \mathcal{V}(\Sigma^{+\infty})$ does in general itself *not* represent a distribution (on traces) but just a (positive) vector in $\mathcal{V}(\Sigma^{+\infty})$. However, if we collect all those traces which start with the same state s then we obtain a distribution over traces, i.e. $\sum \{p \mid \langle \sigma, p \rangle \text{ with } \sigma = s \dots\} = 1$.

It would be possible to formulate the MTS also as a map which expresses the dependency on the initial state explicitly and returns directly distributions over traces, i.e. $\llbracket \cdot \rrbracket_{MTS} : \mathbf{Stmt} \rightarrow \Sigma \rightarrow \mathcal{D}(\Sigma^{+\infty}) \subseteq \mathcal{V}(\Sigma^{+\infty})$, in which case $\llbracket S \rrbracket_{MTS}(s)$ would simply represent a distribution over traces. However, our aim is to stay as close as possible to the formulation in [7], which is based on the typing $\llbracket \cdot \rrbracket_{MTS} : \mathbf{Stmt} \rightarrow \Omega \rightarrow \mathcal{P}(\Sigma^{+\infty})$ rather than, for example, $\llbracket \cdot \rrbracket_{MTS} : \mathbf{Stmt} \rightarrow \Omega \rightarrow \Sigma \rightarrow \mathcal{P}(\Sigma^{+\infty})$.

Example 7. In order to illustrate the basic construction of the MTS we consider the following program in both its unlabelled and labelled version, with $x, y \in \{0, 1\}$:

<pre> if ($y < 1$) then $x \text{ ?= } \{\langle 0, p \rangle, \langle 1, 1 - p \rangle\}$ else $x := 0$ fi; if ($x < 1$) then $y \text{ ?= } \{\langle 0, q \rangle, \langle 1, 1 - q \rangle\}$ else $y := 0$ fi </pre>	<pre> if [$(y < 1)$]¹ then [$x \text{ ?= } \{\langle 0, p \rangle, \langle 1, 1 - p \rangle\}$]² else [$x := 0$]³ fi; if [$(x < 1)$]⁴ then [$y \text{ ?= } \{\langle 0, q \rangle, \langle 1, 1 - q \rangle\}$]⁵ else [$y := 0$]⁶ fi </pre>
---	---

In this example we have no loops or recursions, so we know that we will need (at most) two “coin flips”. Thus, the space of scenarios Ω is defined via the two choices, one for x and one for y , i.e. as $\Sigma = \{x \mapsto 0, x \mapsto 1\} \times \{y \mapsto 0, y \mapsto 1\}$, which we will denote by $\Sigma = \{[00], [01], [10], [11]\}$ with $[00]$ the state $x \mapsto 0, y \mapsto 0$, etc. Following the reformulation of the MTS we have:

$$\begin{aligned}
\llbracket x := 0 \rrbracket_{MTS} &= \\
&= \{ \langle [00][00], 1 \rangle, \langle [01][01], 1 \rangle, \langle [10][00], 1 \rangle, \langle [11][01], 1 \rangle \} \\
\llbracket y := 0 \rrbracket_{MTS} &= \\
&= \{ \langle [00][00], 1 \rangle, \langle [01][00], 1 \rangle, \langle [10][10], 1 \rangle, \langle [11][10], 1 \rangle \} \\
\llbracket x \text{ ?= } \{\langle 0, p \rangle, \langle 1, 1 - p \rangle\} \rrbracket_{MTS} &= \\
&= \{ \langle [00][00], p \rangle, \langle [01][01], p \rangle, \langle [10][00], p \rangle, \langle [11][01], p \rangle, \\
&\quad \langle [00][10], 1 - p \rangle, \langle [01][11], 1 - p \rangle, \langle [10][10], 1 - p \rangle, \langle [11][11], 1 - p \rangle \} \\
\llbracket y \text{ ?= } \{\langle 0, q \rangle, \langle 1, 1 - q \rangle\} \rrbracket_{MTS} &= \\
&= \{ \langle [00][00], q \rangle, \langle [01][00], q \rangle, \langle [10][10], q \rangle, \langle [11][10], q \rangle, \\
&\quad \langle [00][01], 1 - q \rangle, \langle [01][01], 1 - q \rangle, \langle [10][11], 1 - q \rangle, \langle [11][11], 1 - q \rangle \}
\end{aligned}$$

With these sets of weighted traces we can now construct the MTS for the two **if** statements:

$$\begin{aligned}
\llbracket \text{if } (y < 1) \text{ then } x \text{ ?= } \{\langle 0, p \rangle, \langle 1, 1 - p \rangle\} \text{ else } x := 0 \text{ fi} \rrbracket_{MTS} &= \\
&= \{ \langle [00][00][00], p \rangle, \langle [10][10][00], p \rangle, \langle [00][00][10], 1 - p \rangle, \\
&\quad \langle [10][10][10], 1 - p \rangle, \langle [01][01][01], 1 \rangle, \langle [11][11][01], 1 \rangle \} \\
\llbracket \text{if } (x < 1) \text{ then } y \text{ ?= } \{\langle 0, q \rangle, \langle 1, 1 - q \rangle\} \text{ else } y := 0 \text{ fi} \rrbracket_{MTS} &= \\
&= \{ \langle [00][00][00], q \rangle, \langle [01][01][00], q \rangle, \langle [00][00][01], 1 - q \rangle, \\
&\quad \langle [01][01][01], 1 - q \rangle, \langle [10][10][10], 1 \rangle, \langle [11][11][10], 1 \rangle \}
\end{aligned}$$

Note that some traces which we constructed for the branches disappear because when we apply the operator “;” the last state of the first (one step) trace (representing the test) and the first state of the continuation (in one of the two branches) do not match.

The traces for the whole program are then given by:

$$\begin{aligned} \llbracket P \rrbracket_{MTS} = \{ & \langle [00][00][00], p \rangle; \langle [00][00][00], q \rangle, \langle [00][00][00], p \rangle; \langle [00][00][01], 1 - q \rangle, \\ & \langle [10][10][00], p \rangle; \langle [00][00][00], q \rangle, \langle [10][10][00], p \rangle; \langle [00][00][01], 1 - q \rangle, \\ & \langle [00][00][10], 1 - p \rangle; \langle [10][10][10], 1 \rangle, \langle [10][10][10], 1 - p \rangle; \langle [10][10][10], 1 \rangle, \\ & \langle [01][01][01], 1 \rangle; \langle [01][01][00], q \rangle, \langle [01][01][01], 1 \rangle; \langle [01][01][01], 1 - q \rangle, \\ & \langle [11][11][01], 1 \rangle; \langle [01][01][00], q \rangle, \langle [11][11][01], 1 \rangle; \langle [01][01][01], 1 - q \rangle \}, \end{aligned}$$

where again the matching condition eliminates a number of possible traces. Finally we get:

$$\begin{aligned} \llbracket P \rrbracket_{MTS} = \{ & \langle [00][00][00][00][00], pq \rangle, \langle [00][00][00][00][01], p(1 - q) \rangle, \\ & \langle [10][10][00][00][00], pq \rangle, \langle [10][10][00][00][01], p(1 - q) \rangle, \\ & \langle [00][00][10][10][10], 1 - p \rangle, \langle [10][10][10][10][10], 1 - p \rangle, \\ & \langle [01][01][01][01][00], q \rangle, \langle [01][01][01][01][01], 1 - q \rangle, \\ & \langle [11][11][01][01][00], q \rangle, \langle [11][11][01][01][01], 1 - q \rangle \}. \end{aligned}$$

Here we have three possible traces starting with the initial state $s = [00]$ or $s = [10]$ but only two for $s = [01]$ and $[11]$. We also observe that the probabilities associated to the traces starting with each of the four initial states sum up to one, e.g. for $s = [00]$ we have the probabilities $(pq) + (p - pq) + (1 - p) = 1$.

In this presentation of the MTS the states only record the values of the variables but not the current label (or program counter). This makes it possible to obtain the same trace for completely different executions of the program. To keep track of the control flow through the program, its labelled version allows to record in the labels the information about the configurations executed and not just the states. For the labelled version of the program we would then replace a trace like $[00][00][00][00][00]$ by $\langle [00], 1 \rangle \langle [00], 2 \rangle \langle [00], 4 \rangle \langle [00], 5 \rangle \langle [00], \ell^* \rangle$ with ℓ^* the final label indicating termination.

3 Probabilistic Vs Classical Abstract Interpretation

Abstract Interpretation (AI) is a well known mathematical theory at the base of a number of static analysis techniques [4]. Because of the need to consider computable domains for performing the analysis of program properties, abstraction and approximation are essential features of any static analysis technique. The theory of AI establishes when the approximation is such that an analysis can be safely performed on an abstract rather than the concrete domain of computation. More precisely, the correctness of an abstract semantics is guaranteed by ensuring that a pair of functions α and γ can be defined which form a *Galois connection* between two lattices \mathcal{C} and \mathcal{D} representing concrete and abstract properties. This classical theory originally introduced for (non-)deterministic programs can be extended so as to include the treatment of probabilistic programs by considering the appropriate (abstract and concrete) domains as recently shown in [7] (see also [37]).

Though the approximations allowed by the AI theory will always be safe, they might also be quite unrealistic, addressing a *worst case* scenario rather than the *average case* [38]. This latter is typically the aim of a probabilistic analysis which is therefore hardly correct in the classical sense of the AI theory. However, although such an average case analysis is not guaranteed to ‘err on the safe side’, we can still define it so as to reduce the error margin. In order to provide a mathematical framework for probabilistic analysis, we have previously introduced in [5,6], a theory of linear operators on Hilbert spaces (i.e. here just finite dimensional spaces as discussed before) where the notion of approximation is characterised in terms of *least square approximation*, which we have called *Probabilistic Abstract Interpretation* (PAI).

The PAI approach is based, as in the classical case, on a concrete and abstract domain \mathcal{C} and \mathcal{D} – except that \mathcal{C} and \mathcal{D} are now vector spaces instead of lattices. We assume that the pair of abstraction and concretisation function $\mathbf{A} : \mathcal{C} \rightarrow \mathcal{D}$ and $\mathbf{G} : \mathcal{D} \rightarrow \mathcal{C}$ are again structure preserving, i.e. in our setting they are (bounded) linear maps represented by matrices \mathbf{A} and \mathbf{G} . Finally, we replace the notion of a Galois connection by the notion of *Moore-Penrose pseudo-inverse* [8,10].

Definition 1. *Let \mathcal{C} and \mathcal{D} be two finite dimensional vector spaces, and let $\mathbf{A} : \mathcal{C} \rightarrow \mathcal{D}$ be a linear map between them. The linear map $\mathbf{A}^\dagger = \mathbf{G} : \mathcal{D} \rightarrow \mathcal{C}$ is the Moore-Penrose pseudo-inverse of \mathbf{A} iff*

$$\mathbf{A} \circ \mathbf{G} = \mathbf{P}_A \quad \text{and} \quad \mathbf{G} \circ \mathbf{A} = \mathbf{P}_G$$

where \mathbf{P}_A and \mathbf{P}_G denote orthogonal projections (i.e. $\mathbf{P}_A^* = \mathbf{P}_A = \mathbf{P}_A^2$ and $\mathbf{P}_G^* = \mathbf{P}_G = \mathbf{P}_G^2$ where $*$ denotes the adjoint [33, Ch 10]) onto the ranges of \mathbf{A} and \mathbf{G} .

Alternatively, if \mathbf{A} is Moore-Penrose invertible (and all finite dimensional operators or matrices are), its Moore-Penrose pseudo-inverse, \mathbf{A}^\dagger satisfies the following:

- (i) $\mathbf{A}\mathbf{A}^\dagger\mathbf{A} = \mathbf{A}$,
- (ii) $\mathbf{A}^\dagger\mathbf{A}\mathbf{A}^\dagger = \mathbf{A}^\dagger$,
- (iii) $(\mathbf{A}\mathbf{A}^\dagger)^* = \mathbf{A}\mathbf{A}^\dagger$,
- (iv) $(\mathbf{A}^\dagger\mathbf{A})^* = \mathbf{A}^\dagger\mathbf{A}$.

It is instructive to compare these equations with the classical setting. For example, a Galois connection (α, γ) satisfies the properties $\alpha \circ \gamma \circ \alpha = \alpha$ and $\gamma \circ \alpha \circ \gamma = \gamma$ which are similar to conditions (i) and (ii) in Definition 1. Moreover, we also have in a similar way as in the AI setting that \mathbf{A} and \mathbf{A}^\dagger determine each other uniquely, i.e. $(\mathbf{A}^\dagger)^\dagger = \mathbf{A}$ (cf. e.g. [10]).

The Moore-Penrose pseudo-inverse allows us to construct the closest (i.e. least square) approximation $\mathbf{T}^\# : \mathcal{D} \rightarrow \mathcal{D}$ of a concrete semantics $\mathbf{T} : \mathcal{C} \rightarrow \mathcal{C}$ as:

$$\mathbf{T}^\# = \mathbf{G} \cdot \mathbf{T} \cdot \mathbf{A} = \mathbf{A}^\dagger \cdot \mathbf{T} \cdot \mathbf{A} = \mathbf{A} \circ \mathbf{T} \circ \mathbf{G}.$$

In [5] we show how we can transform a Probabilistic Abstract Interpretation into a classical Abstract Interpretation by forgetting the concrete values of probabilities and only considering the support set of a distribution as the set of “possibilities”. One can also lift (in a non-unique way) a classical Abstract Interpretation to a Probabilistic Abstract Interpretation (e.g. by using uniform distributions). This method is conceptually equivalent to the probabilistic version of Abstract Interpretation presented in [7], although the result does not refer explicitly to the Maximal Trace Semantics. However, AI and PAI are not equivalent in terms of the analyses that they support. Besides the relaxation of the safety constraint for the analysis results, PAI is also a suitable mathematical framework for *testing*, as we will show in Sect. 5.

4 Comparison of Probabilistic Semantics

For the language **pWhile**, the Kozen semantics describes the I/O behaviour of programs, the LOS semantics gives the generator for a step-wise execution of the program (as a DTMC), and the MTS determines the possible traces and their corresponding probabilities (inherited from the scenarios of the probability space). In this section we will discuss in some detail the relationship between them with the aim of clarifying their different role in the static analysis of programs.

Kozen’s Semantics and LOS. One important difference between the LOS and Kozen’s semantics (Semantics 2 in [12]) is the use of labels (as a kind of program counter) to model the computational steps.

As already mentioned, in Kozen’s semantics all non-terminating executions are treated equally, i.e. have a trivial or zero semantics. Another difference is that Kozen’s semantics is based on a state space $\mathcal{V}(\mathbf{Value}^n)$ as opposed to the LOS state space $\mathcal{V}(\mathbf{Value})^{\otimes n}$ which allows for an independent treatment of each variable. In general, the tensor construction of the LOS allows for a kind of ‘compositional’ program analysis where the various syntactic components of a program can be analysed individually, which is not possible with Kozen’s semantics.

In [28] we have shown that Kozen’s operator $\llbracket P \rrbracket_{KFS}$ is an abstraction of a limit of iterations of the LOS semantics $\llbracket P \rrbracket_{LOS}$. This abstraction is defined by the PAI operator which “forgets” about the computational state at all labels except ℓ :

$$\mathbf{A}_\ell = \mathbf{I} \otimes \dots \otimes \mathbf{I} \otimes e_\ell,$$

where e_ℓ is a unit or base vector in $\mathcal{V}(\mathbf{Label})$ corresponding to label $\ell \in \mathbf{Label}$, i.e. $e_\ell = (0, 0, \dots, 0, 1, 0, \dots, 0)$ with only one non-zero entry for the coordinate ℓ . This can also be seen as $1 \times |\mathbf{Label}|$ matrix. This operation keeps all the information about the state, i.e. values of the variables, but only when the execution is in label ℓ . If we take $\ell = \ell^*$, i.e. the terminal looping state in the semantics of a program, then this gives the probabilities of the values of all variables for those computations which have already reached the end. So for any initial (classical)

state s_0 and initial label $\ell = 0$ we can obtain the computational state in the final label ℓ^* by iteration. The following propositions hold (cf. [28]):

Proposition 1. *Given a **pWhile** program P and initial state s_0 in $\mathcal{V}(\mathbf{Value})^{\otimes n}$, then $(s_0 \otimes e_0)[P]_{LOS}^t \mathbf{A}_{\ell^*}$ corresponds to the distributions over all states on which P terminates in t or fewer computational steps.*

This covers all finite computations of t steps or fewer. In order to get the I/O behaviour for all terminating computations, i.e. the Kozen semantics, we need just to consider the limit of all computations of any length:

Proposition 2. *Given a **pWhile** program P and initial state s_0 in $\mathcal{V}(\mathbf{Value})^{\otimes n}$, let $[P]_{KFS}$ be Kozen's semantics of P and $[P]_{LOS}$ the DTMC generator for P . Then*

$$(s_0 \otimes e_0) \left(\lim_{t \rightarrow \infty} [P]_{LOS}^t \right) \mathbf{A}_{\ell^*} = s_0 [P]_{KFS}.$$

Maximal Trace Semantics and LOS. The probabilistic semantics in [7] is a classical abstraction of the probabilistic MTS corresponding to a “strongest post-condition semantics”. This is in effect an operator semantics which maps input distributions into some output distributions (cf. formula (2) in [7, 7.4]). A common interpretation of the claims made in Sect. 7.3 of [7] is that the LOS is just an abstraction of the probabilistic MTS. We show here that this interpretation is incorrect.

In order to investigate the relationship between LOS and MTS in more detail we will look at a concrete construction of probabilistic traces by means of the LOS. It is somewhat unclear if the MTS in [7] should be based on $\Sigma = \mathbf{State}$ or $\Sigma = \mathbf{Conf}$. In the first case it is straightforward to see that the LOS actually contains more information than an MTS based only on state information. We will thus consider the MTS based on $\Sigma = \mathbf{Conf}$, i.e. the reformulation of the probabilistic MTS as an element in $\mathcal{V}(\mathbf{Conf}^{+\infty})$, which associates with every possible trace a probability that this is indeed the trace which will be executed during the program run. We will then relate this set of ‘weighted’ traces to the LOS as an operator on $\mathcal{V}(\mathbf{Conf})$ where we also provide the initial distribution $s_o \otimes e_0 = \rho_0 \in \mathcal{V}(\mathbf{Conf})$.

The LOS allows for the construction of a *sequence of distributions* over states (fronts): Given an initial state we can calculate for every t the probabilities of reaching any state after t steps by applying the LOS operator t times to the initial state. The MTS does not construct fronts but rather a *distribution over sequences* (traces): given an initial state we can calculate the probabilities of all the execution traces starting from that initial state. The two notions are thus somewhat orthogonal. However, they turn out to be equivalent for languages that, like **pWhile**, can be modelled via a DTMC. This is because DTMC’s abstract from the history of a computation as only the current configuration determines the probabilities of the successor configurations. Transition probabilities are exactly what is specified in the generator matrix of the DTMC and is all one needs to reconstruct both the fronts and the computational traces with their probabilities.

Instantiated for purely probabilistic languages the classical abstraction given by formula (2) in [7, 7.4] is an operator from distributions over traces to distribution transformers (for a fixed initial configuration s), i.e.

$$\alpha_s : \mathcal{V}(\mathbf{Conf}^{+\infty}) \rightarrow \mathcal{L}(\mathcal{V}(\mathbf{Conf}))$$

rather than $(\Omega \rightarrow \mathcal{P}(\mathbf{Conf}^{+\infty})) \rightarrow (\mathcal{V}(\mathbf{Conf}) \rightarrow \mathcal{V}(\mathbf{Conf}))$ as in [7, 7.4]. In this purely probabilistic case, the abstraction map becomes:

$$((\alpha_s(\{\langle p, X \rangle\})(\delta))(s')) = \sum_{s \in \Sigma} \{\delta(s) \cdot p \mid \text{for } s\sigma s' \in X^+\}.$$

In other words, we associate to every distribution over traces $\{\langle p, X \rangle\}$ a linear operator $(\alpha_s(\{\langle p, X \rangle\})) \in \mathcal{L}(\mathcal{V}(\mathbf{Conf}))$. To see how this operator transforms a distribution $\delta \in \mathcal{V}(\mathbf{Conf})$ into another distribution $\alpha_s(\{\langle p, X \rangle\})(\delta) = \delta' \in \mathcal{V}(\mathbf{Conf})$ we describe the probability of every configuration $s' \in \mathbf{Conf}$ in the new distribution δ' . This is the sum of all products of the probabilities associated with all the traces which, starting from any s , reach s' in finitely many steps and the probability $\delta(s)$ that we start indeed with s . The probability $\delta'(s')$ is the probability that we terminate with s' . Therefore this abstraction gives the Kozen I/O semantics. However, it does not give the LOS which instead would require a classical abstraction of the form

$$((\bar{\alpha}_s(\{\langle p, X \rangle\})(\delta))(s')) = \sum_{s \in \Sigma} \{\delta(s) \cdot p \mid \text{if } ss' \dots \in X\}$$

i.e. an operator that collects the probabilities that we reach s' in one step rather than eventually. Note that this abstraction does not require that s' is a terminating state. The question is now whether $\bar{\alpha}_s$ is indeed an abstraction or not. If we consider the dimension of the spaces involved, the answer is positive as there is obviously a loss of information when considering the space $\mathcal{L}(\mathcal{V}(\mathbf{Conf}))$ of $n \times n$ matrices (for \mathbf{Conf} with n states) with dimension n^2 in place of the space of distributions $\mathcal{V}(\mathbf{Conf})^{\otimes t}$ (on traces of finite length t) whose dimension is n^t . However, due to the memory-less property of DTMC, we only need to consider traces of length 2 (i.e. transition steps) and thus a space $\mathcal{V}(\mathbf{Conf}) \otimes \mathcal{V}(\mathbf{Conf})$ whose dimension is n^2 . Thus, no information is lost and the abstraction is not really an abstraction but only a recasting of the MTS. If the MTS is the most concrete semantics (in the sense of [36]) then so is the LOS. In fact, we can show the following proposition.

Proposition 3. *Given a pWhile program P , then the LOS $\llbracket P \rrbracket_{LOS}$ and the MTS $\llbracket P \rrbracket_{MTS}$ are equivalent, i.e. it is possible to construct either semantics from the other one.*

It is straightforward to construct the LOS operator out of the MTS by considering for all initial configurations (i.e. point-distributions) the single step traces (or single step trace-prefixes) starting from each initial configuration. In fact,

the probability associated to these traces is exactly the transition probability recorded in the DTMC generator, i.e. the LOS – this is indeed what the map $\bar{\alpha}_s$ above achieves. On the other hand, the probability associated with a trace $s_{i_1} s_{i_2} s_{i_3} \dots$ is the product of the transition probabilities $(\llbracket P \rrbracket_{LOS})_{i_1 i_2}$, $(\llbracket P \rrbracket_{LOS})_{i_2 i_3}$ etc. – i.e. $\prod_j (\llbracket P \rrbracket_{LOS})_{i_j i_{j+1}}$ – times the probabilities given by the initial distribution $\delta(s_{i_1})$.

As already mentioned, these constructions require that the semantics of **pWhile** is modelled by a *homogeneous* DTMC, i.e. that the transition probabilities from one configuration to another one do not change over time. This and the memory-less property of DTMC’s seems to be a reasonable requirement for a programming language.

5 Statistical Analysis of Probabilistic Programs via PAI

Probabilistic semantics provides the basis for the static analysis of probabilistic programs. While both the AI and the PAI framework allow us to use traces as a basis for constructing more abstract semantics, there is an important difference between the two frameworks. In the AI setting these traces are assumed to be *ideal* traces, i.e. traces that are actually obtained when a program is executed. In the PAI setting – similar to the situation in statistical analysis, learning etc. – we can attempt to utilise not just ideal traces but also experimentally *observed*, maybe *corrupted*, i.e. *distorted* by noise, traces in order to reconstruct the most plausible underlying abstract semantics.

In this section we show an approach where the probabilistic information about the program executions is inferred by *observing* some sample runs. This establishes a link between static program analysis and testing and demonstrates the use of PAI to calculate best estimates for program’s properties in a way similar to the so-called *linear statistical model* or *linear regression* method.

The approach we are going to present is based on the idea of identifying observations with a linear combination of a set of random variables x_i , whose weights are chosen with the method of least squares so as to minimise the distance from the observations and the actual model expressing the program behaviour. Thus the framework of Probabilistic Abstract Interpretation is particularly appropriate as a base of this approach.

5.1 The Linear Statistical Model

In several contexts it is often useful to predict or estimate a variable β (or a vector of variables), given that we have the opportunity to observe variables y_1, y_2, \dots, y_n that somehow (statistically) depend on β . This is a very important statistical problem which is typically faced by using so-called linear regression analysis, also known as linear statistical model (cf e.g. [11], [10, Section 8.3] or [8, Section 6.4]). This widely used statistical technique applies to situations such as the one mentioned above, where a random vector y depends *linearly* on a vector of parameters β , i.e. (using post-multiplication)

$$y = \beta \mathbf{X} + \varepsilon, \tag{1}$$

where y represents some measurement results, the parameters β are unknown, the matrix \mathbf{X} is the *design matrix*, and ε is a random vector representing the errors of observing y . This error is conventionally assumed to have expected value equal to zero and some further statistical conditions regarding its variance and co-variance are typically imposed. These requirements mean that there is no underlying or systematic reason for the distortions ε and this is only due to random noise.

The role of least square approximations and the Moore-Penrose pseudo-inverse in this context is of particular relevance for the well-known Gauss-Markov theorem (cf. [10, Section 8.3, Thm. 1]).

Theorem 1 (Gauss-Markov). *Consider the linear model $y = \beta\mathbf{X} + \varepsilon$ with \mathbf{X} of full column rank and ε fulfilling the conditions in [10, Section 8.3]. Then the Best Linear Unbiased Estimator (BLUE) is given by*

$$\hat{\beta} = y\mathbf{X}^\dagger.$$

In its simplest version, the Gauss-Markov theorem thus asserts that the best estimate $\hat{\beta}$ of the unknown parameters β can be obtained from some experimentally observed y by calculating $y\mathbf{X}^\dagger$, i.e. via the Moore-Penrose pseudo-inverse of the design matrix \mathbf{X} , cf. [10, Section 8.3, eqn (35)].

5.2 Application to Security Analysis

We discuss the relevance of the reconstruction of unknown parameters or properties of a system in the field of computer security by presenting a simplified version of the well-known Kocher’s attack on crypto-protocols [39].

Modular exponentiation is a basic operation for computing the private key in crypto-systems using the Diffie-Hellman or the RSA protocols. In [39], it is shown that by carefully measuring the time required to perform such an operation, an attacker may be able to find the Diffie-Hellman exponents or factor the RSA keys and break the crypto-systems.

The crucial point is the estimation of a single bit b in the secret key k . Since modular exponentiation takes very different execution times depending on the value of a certain bit b being 0 or 1, what the attacker needs are good estimate of these execution times in order to deduce the value of each bit of the key. Thus, linear statistical models play a crucial role in the analysis of security. We show how the problem of the timing attacks can be described as a statistical analysis problem, by using as an example a simplified implementation of the RSA exponentiation algorithm. This will also highlight the relationship between PAI and linear regression.

Suppose that t_0 is the time it takes to perform multiplication in the modular exponentiation procedure if a single bit b of the cryptographic key k is $b = 0$ and t_1 if $b = 1$. We thus need to consider two possible DTMC models, one for the case $b = 0$ and one for $b = 1$. In realistic situations we also need to take into account the noise due for example to the fact that the physical device we

observe is also involved in other tasks/threads such as network communication. The aim is to guess correctly which of the two models is actually being executed, i.e. the value of b , by observing the (possibly distorted) running time. We can also set the vector β to represent the strength/weights/probabilities that in a given model we have $b = 0$ or $b = 1$, respectively. More concretely, we can set the vector $\beta_0 = (1, 0)$ to represent the models of the system where the bit b of the key is $b = 0$ and $\beta_1 = (0, 1)$ to the key with $b = 1$. We can now define a linear statistical model by constructing a design matrix \mathbf{X} (in the PAI sense a concretisation operator), which maps a model (element in the abstract domain) onto its timing behaviour (element in the concrete domain). As an example, we can consider the situation where we can observe ten possible execution times t_i that we enumerate and use as column indices for \mathbf{X} . Suppose that t_0 corresponds to the 3rd and t_1 to the 7th column in this enumeration. In this case we obtain a design matrix of the form:

$$\mathbf{X} = \begin{pmatrix} 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \end{pmatrix},$$

and we can calculate

$$\beta_0 \mathbf{X} = (1, 0) \cdot \mathbf{X} = (0, 0, 1, 0, 0, 0, 0, 0, 0, 0)$$

which tells us that for $b = 0$ the chances of observing any other time signature than t_0 is zero, and that t_0 will definitively be observed. A similar calculation can be done for $\beta_1 = (0, 1)$.

If we begin instead by observing the time behaviour, i.e. if we test the program and obtain, for example, an (undistorted) observation vector of the form

$$y = (0, 0, 1, 0, 0, 0, 0, 0, 0, 0),$$

then, by calculating $y\mathbf{X}^\dagger = (1, 0)$ we will get that b is definitely 0. If we now add a (Gaussian) error to our experiment then the observed times, corresponding to an estimate y would perhaps be something like (cf. Figures 1 and 2 in [39]):

$$\hat{y} = (0.1, 0.2, 0.7, 0.2, 0.1, 0, 0, 0, 0, 0),$$

because, for example, in 10 measurements we have observed once the first possible time, twice the second, etc. The estimation based on these observations leads to a guess of the weights of the parameters in β that we calculate as $\hat{y}\mathbf{X}^\dagger = (0.7, 0)$. This result reflects the fact that it is very likely that the value of bit b is 0 as we have observed, although with some errors, a time behaviour where the times cluster around the value t_0 .

5.3 Abstraction and Linear Regression


Statistics can be used in static analysis in all those cases where we have some observations at hand and we want to use them in order to improve the precision

of the analysis. To this purpose the theory of linear regression provides us with a useful means to determine a best estimate of the model underlying those observations, e.g. the DTMC generator that with highest probability produces the traces that we observe.

Note that classical abstract interpretation cannot be used in this scenario even in its probabilistic re-formulation as given e.g. in [7]; this is because the safety constraint at the base of the framework does not permit the consideration of expectation values in the analysis result, as these would not guarantee the correctness of the analysis (cf. Section 3).

In the setting of linear statistical models, the concretisation operator \mathbf{G} of the PAI framework corresponds to a mapping from an abstract domain consisting of all possible DTMC models for the observed program to all possible observable traces corresponding to the different runnings of the program. Thus, \mathbf{G} plays the role of the design matrix of the statistical model. If y is a vector defining the probabilities of certain traces according to some observations and β represents a parameterised DTMC model, then we can use the linear statistical Eq. (1) in its simplest instance, i.e. with $\varepsilon = 0$, $y \in \mathbb{R}^n$, $\mathbf{X} \in \mathbb{R}^{n \times p}$ and $\beta \in \mathbb{R}^p$, in order to obtain the best estimate of the concrete DTMC model by $\hat{\beta} = y\mathbf{X}^\dagger$.

Example 8. Consider the following simple examples of DTMC's:




with generator matrices

$$\mathbf{T}_{0,1} = \begin{pmatrix} 0 & 1 \\ 0 & 1 \end{pmatrix} \quad \mathbf{T}_{\frac{1}{2},1} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ 0 & 1 \end{pmatrix} \quad \mathbf{T}_{0,\frac{1}{2}} = \begin{pmatrix} 0 & 1 \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix} \quad \mathbf{T}_{\frac{1}{2},\frac{1}{2}} = \begin{pmatrix} \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix}.$$

Clearly with $\mathbf{T}_{\frac{1}{2},\frac{1}{2}}$ we can generate all infinite 0/1 sequences. Note that since these are uncountably many, the probability structure on the maximal trace space will require a measure theoretical treatment. By restricting ourselves to traces of finite length we can however stay within a finite-dimensional setting.

These DTMC's are in essence the processes which, for different values of p and q , describe the core (loop body) of Example 1 in the LOS or Kozen semantics (cf. also $\llbracket S \rrbracket_{KFS}$ in Example 3).

The processes above depend on the parameters p and q in the real interval $[0,1]$ which we can see as the probability to remain in state 0 and state 1, respectively. They are represented by the parametric DTMC



with generator

$$\mathbf{T}_{pq} = \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix}$$

Any property of a program whose behaviour can be described as above will depend on the parameters p and q . Moreover, observing the property may be influenced by some distorted execution of \mathbf{T}_{pq} . Applying the statistical linear model to find best estimates for the parameters p and q corresponds to performing a statistical analysis based on PAI, as shown in the the following example.

Example 9. Consider the DTMC in Example 8 with generator

$$\mathbf{T}_{pq} = \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix}.$$

This system is completely specified when both the values of p and q and the initial state s are specified. Thus we can identify the abstract semantic domain with the set of all pairs of initial states $s \in \{0, 1\}$ and matrices \mathbf{T}_{pq} , i.e.

$$\mathcal{M} = \{\langle s, \mathbf{T}_{pq} \rangle\} = \left\{ \left\langle s, \begin{pmatrix} p & 1-p \\ 1-q & q \end{pmatrix} \right\rangle \right\}$$

or equivalently with the set of triples $\mathcal{M} = \{\langle s, p, q \rangle \mid s \in \{0, 1\}, p, q \in [0, 1]\}$.

Note that this parametric DTMC generator encodes the same information as the set of all parametric traces starting from any initial state (cf. Section 4). In order to apply PAI we consider the distributions over \mathcal{M} , i.e. the space $\mathcal{D} = \mathcal{V}(\mathcal{M})$ of all normalised, positive elements in the vector space over \mathcal{M} .

The concrete computational space consists of the set of all sequences of 0 and 1 in $\mathcal{T} = \{0, 1\}^{+\infty}$, representing the execution traces resulting from fixing actual values of the parameters p and q and the input state. The concrete domain of PAI is therefore the space of distributions on traces $\mathcal{C} = \mathcal{V}(\mathcal{T})$.

Numerical Experiments. Even for the simple example given above, the sets involved are uncountably infinite. In order to be able to compute an analysis of the system in Example 8 we will consider here the simple case where transition probabilities can only assume values in a finite set, i.e. $p, q \in \{p_0, \dots, p_n\}$ and where traces can only be of length t , for a given t . We report below some of the results we obtained from numerical experiments performed using the Octave system [40]. In these experiments we considered $p, q \in \{0, \frac{1}{2}, 1\}$, thus obtaining 9 possible semantics, with possible initial states either 0 or 1. This corresponds to

an abstract domain $\mathcal{D} = \mathcal{V}(\{0, 1\}) \otimes \mathcal{V}(\{0, \frac{1}{2}, 1\}) \otimes \mathcal{V}(\{0, \frac{1}{2}, 1\}) = \mathbb{R}^2 \otimes \mathbb{R}^3 \otimes \mathbb{R}^3 = \mathbb{R}^{18}$.

For different models – i.e. different values of p and q – as well as different noise levels we simulated 10000 executions of the system and observed traces of length $t = 10$. In this setting, concrete domain is therefore $\mathcal{C} = \mathcal{V}(\{0, 1\}^{10}) = \mathcal{V}(\{0, 1\})^{\otimes 10} = (\mathbb{R}^2)^{\otimes 10} = \mathbb{R}^{1024}$, i.e. there are about one thousand possible traces that can be observed.

The concretisation/design matrix $\mathbf{G} : \mathcal{D} \rightarrow \mathcal{C}$ associates to each of the 18 instance models and initial inputs one of the distributions over the 1024 traces, namely the one representing those traces that are obtained in that model. As it is impossible to reproduce here the actual 18×1024 matrix \mathbf{G} (due to its size) we give in Fig. 1 the matrix G for the restricted case of 8 possible traces of 3 steps, with rows representing the possible instance models and columns the possible traces. The entries of this matrix specify the probabilities that a given model (row) generates a certain trace (column). For example, the entry $\mathbf{G}_{33} = \frac{1}{2}$ means that with the third model in the enumeration given above, i.e. for initial state $s = 0$, $p = \frac{1}{2}$ and $q = 0$, we get the third trace, i.e. 010, with probability $\frac{1}{2}$.

In order to calculate the best estimators of the parameters p and q , we computed the Moore-Penrose pseudo-inverse \mathbf{G}^\dagger of \mathbf{G} , which is also reported in Fig. 1 for the restricted case. Intuitively, \mathbf{G}^\dagger gives us the probabilities that when a certain trace is observed this comes from a certain model.

In our experiments we considered systems without distortion, i.e. no error, as well as the cases where a noise of “strength” ε was applied according to a normal distribution (cf. `randn()` in Octave 3.8.0 [40, p391]).

The observations were aggregated to a distribution over all 2^{10} possible traces. The probability associated to each trace σ_i is the ratio between the number of times σ_i was actually observed and the number of experiments we ran (i.e. 10000 in our case). For the undistorted case we denote this distribution vector by y , for $\varepsilon = 0.01$ by y' , for $\varepsilon = 0.1$ by y'' , and for $\varepsilon = 0.25$ by y''' . The initial state was always chosen with probability $\frac{1}{2}$ as state 0 or state 1.

Model $p = 0 = q$: The vectors we obtained in the case when the true model is given by $p = 0 = q$ are for the different noise levels:

$$\begin{aligned}
 y\mathbf{G}^\dagger &= (0.50 \ 0.50 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y'\mathbf{G}^\dagger &= (0.47 \ 0.49 \ 0.02 \ 0.01 \ 0 \ 0 \ 0.01 \ 0.03 \ -0.02 \ -0.02 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y''\mathbf{G}^\dagger &= (0.33 \ 0.34 \ 0.17 \ 0.11 \ 0 \ 0 \ 0.11 \ 0.18 \ -0.12 \ -0.12 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y'''\mathbf{G}^\dagger &= (0.18 \ 0.17 \ 0.28 \ 0.18 \ 0 \ 0 \ 0.18 \ 0.26 \ -0.13 \ -0.12 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0)
 \end{aligned}$$

Model $p = \frac{1}{2} = q$: The same observations for the case that $p = \frac{1}{2} = q$ gave us:

$$\begin{aligned}
 y\mathbf{G}^\dagger &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0.51 \ 0.50 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y'\mathbf{G}^\dagger &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0.51 \ 0.49 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y''\mathbf{G}^\dagger &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0.51 \ 0.49 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0) \\
 y'''\mathbf{G}^\dagger &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0.50 \ 0.50 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0 \ 0)
 \end{aligned}$$




Fig. 1. Relating models (s, p, q) and traces of length 3

Model $p = 0, q = \frac{1}{2}$: Finally, for the case of an underlying model with $p = 0$ and $q = \frac{1}{2}$ we obtained:

$$\begin{aligned}
 y\mathbf{G}^\dagger &= (0\ 0\ 0\ 0\ 0\ 0\ 0.50\ 0.49\ 0\ 0.01\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0) \\
 y'\mathbf{G}^\dagger &= (0\ 0\ 0\ 0\ 0\ 0\ 0.49\ 0.50\ 0.01\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0) \\
 y''\mathbf{G}^\dagger &= (0\ 0\ 0\ 0\ 0\ 0\ 0.43\ 0.43\ 0.07\ 0.06\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0) \\
 y'''\mathbf{G}^\dagger &= (0\ 0\ 0.01\ 0\ 0\ 0\ 0.33\ 0.35\ 0.16\ 0.16\ 0\ 0\ 0\ 0\ 0\ 0\ 0\ 0)
 \end{aligned}$$

These results demonstrate that if we observe the *undisturbed* DTMC in order to obtain experimentally the probabilities y for all possible 2^{10} traces then we can identify the underlying model more or less uniquely. The abstract distribution $\hat{\beta} = y\mathbf{G}^\dagger$ (i.e. when $\varepsilon = 0$) gives an estimate which corresponds precisely to the true parameters p and q and the probability $\frac{1}{2}$ for the initial states 0 and 1 (cf. the enumeration of models in Fig. 1).

For $\varepsilon = 0.01$ we can also identify the unknown system with high probability. However, there are coordinates of $y'\mathbf{G}^\dagger$ which are non-zero although they do not correspond to the actual system. These stem from the fact that y' has non-zero

probability for traces which actually should not be realised but due to the noise distortion are nevertheless observed.

If we increase the error term in the simulation, i.e. for the distortion $\varepsilon = 0.1$ or $\varepsilon = 0.25$, the possibility of a wrong identification of the actual model(s) is (as expected) higher: The weights associated to the actual system tends to decrease further, while other possible models get stronger. If we further increase ε the estimate $\hat{\beta}$ will still be the optimal one (BLUE) but ultimately it will not allow any meaningful identification of the actual system – we will get only (white) noise. We obtained similar results also for other choices of p and q .

6 Conclusions

We have presented a comparison of three different probabilistic semantics: (i) Kozen’s I/O Fixed-Point Semantics, (ii) the Linear Operator Semantics previously introduced by the authors, and (iii) a probabilistic version of the Maximal Trace Semantics. We have argued that Kozen’s semantics can be recovered as an abstract limit from the LOS (cf. [28]) and that the abstraction α_s in [7, Section 7.4] in fact gives Kozen’s semantics (by collecting the information/probability along finite traces in the MTS). We also demonstrated that LOS contains more information than MTS, namely information about the control flow, but that otherwise LOS and MTS are equivalent.

The second part of this paper relates the Probabilistic Abstract Interpretation framework introduced in [5] with the most widely used statistical technique, namely Linear Regression. As already shown in [5], classical Abstract Interpretations can be recovered from a Probabilistic Abstract Interpretation by means of a forgetful functor that restricts probabilistic domains to their support sets. In this paper we have extended the (re)construction of the LOS from the MTS alluded to in [7] – though this involves the “abstraction” $\bar{\alpha}_s$ rather than α_s – to deal also with distorted observations of traces. This provides a bridge between statistics (testing) and static program analysis. Intended application areas include problems in computer security like covert channels and non-interference notions reinterpreted as process equivalence.

Our presentation was restricted to finite state spaces. However a full treatment of the different semantical models is possible though slightly more complex as it involves a deeper study of the underlying measure-theoretic notion (e.g. the σ -algebras generated by trace pre-fixes) as well as topological notions (e.g. Hilbert vs Banach spaces and their operators, weak limits etc., cf. [28]).

Finally, it might be worth pointing out the rich literature on filtering, system identification, Hidden Markov Models (e.g. [41–43]), and related topics which we did not discuss but are clearly related. Our approach to Linear Regression could be considered to be very simple and basic. However, we think it is worth highlighting the relationship between PAI and statistics. Given the role that least square methods – i.e. the Moore-Penrose pseudo-inverse – play in control theory etc. – for example, for the well-known and celebrated technique of Kalman filters [11] – we aim to further explore this field.

References

1. Jones, N.D., Nielson, F.: Abstract interpretation: a semantics-based tool for program analysis. In: Handbook of Logic in Computer Science, pp. 527–636. Clarendon Press, Oxford (1985)
2. Nielson, F.: Strictness analysis and denotational abstract interpretation. *Inf. Comput.* **76**(1), 29–92 (1988)
3. Nielson, F., Nielson, H.R.: Infinitary control flow analysis: a collecting semantics for closure analysis. In: Proceedings of POPL 1997, pp. 332–345 (1997)
4. Nielson, F., Nielson, H.R., Hankin, C.: Principles of Program Analysis. Springer, Heidelberg (1999)
5. Di Pierro, A., Wiklicky, H.: Concurrent constraint programming: towards probabilistic abstract interpretation. In: Proceedings of PPDP 2000, pp. 127–138 (2000)
6. Di Pierro, A., Wiklicky, H.: Measuring the precision of abstract interpretations. In: Lau, K.-K. (ed.) LOPSTR 2000. LNCS, vol. 2042, pp. 147–164. Springer, Heidelberg (2001)
7. Cousot, P., Monerau, M.: Probabilistic abstract interpretation. In: Seidl, H. (ed.) Programming Languages and Systems. LNCS, vol. 7211, pp. 169–193. Springer, Heidelberg (2012)
8. Campbell, S.L., Meyer, C.D.: Generalized Inverses of Linear Transformations. Pitman - Dover, London (1979)
9. Deutsch, F.: Best Approximation in Inner-Product Spaces. Springer, New York (2001)
10. Ben-Israel, A., Greville, T.N.E.: Generalized Inverses - Theory and Applications. CMS Books in Mathematics, 2nd edn. Springer, New York (2003)
11. Albert, A.: Regression and the Moore-Penrose Pseudoinverse. Mathematics in Science and Engineering. Elsevier, New York (1972)
12. Kozen, D.: Semantics of probabilistic programs. *J. Comput. Syst. Sci.* **22**(3), 328–350 (1981)
13. Di Pierro, A., Sotin, P., Wiklicky, H.: Relational analysis and precision via probabilistic abstract interpretation. In: Proceedings of QAPL 2008. vol. 220(3) of ENTCS, pp. 23–42. Elsevier (2008)
14. Di Pierro, A., Hankin, C., Wiklicky, H.: Probabilistic semantics and program analysis. Formal Methods for Quantitative Aspects of Programming Languages. LNCS, vol. 6154, pp. 1–42. Springer, Heidelberg (2010)
15. Di Pierro, A., Hankin, C., Wiklicky, H.: Probabilistic lambda calculus and quantitative program analysis. *J. Logic Comput.* **15**(2), 159–179 (2005)
16. Ramsey, N., Pfeffer, A.: Stochastic lambda calculus and monads of probability distributions. *ACM SIGPLAN Notices* **37**(1), 154–165 (2002)
17. Pfeffer, A.: Practical Probabilistic Programming. Manning, Shelter Island (2015)
18. Di Pierro, A., Hankin, C., Wiklicky, H.: Probabilistic linda-based coordination languages. In: de Boer, F.S., Bonsangue, M.M., Graf, S., de Roever, W.-P. (eds.) FMCO 2004. LNCS, vol. 3657, pp. 120–140. Springer, Heidelberg (2005)
19. Priami, C.: Stochastic π -calculus. *Comput. J.* **38**(7), 578–589 (1995)
20. Hillston, J.: A Compositional Approach to Performance Modelling. Cambridge University Press, Cambridge (1996)
21. Di Pierro, A., Hankin, C., Wiklicky, H.: A systematic approach to probabilistic pointer analysis. In: Shao, Z. (ed.) APLAS 2007. LNCS, vol. 4807, pp. 335–350. Springer, Heidelberg (2007)

22. Jones, C., Plotkin, G.D.: A probabilistic powerdomain of evaluations. In: Proceedings of LICS 1989, pp. 186–195. IEEE (1989)
23. Jones, C.: Probabilistic non-determinism. Ph.D. thesis, University of Edinburgh (1989)
24. Jung, A., Tix, R.: The troublesome probabilistic powerdomain. *ENTCS* **13**, 70–91 (1998)
25. Morgan, C., McIver, A., Seidel, K.: Probabilistic predicate transformers. *ACM Trans. Program. Lang. Syst.* **18**(3), 325–353 (1996)
26. Gretz, F., Katoen, J.P., McIver, A.: Operational versus weakest pre-expectation semantics for the probabilistic guarded command language. *Perform. Eval.* **73**, 110–132 (2014)
27. Park, S., Pfenning, F., Thrun, S.: A probabilistic language based upon sampling functions. In: Proceedings of POPL 2005, 171–182 (2005)
28. Di Pierro, A., Wiklicky, H.: Semantics of probabilistic programs: a weak limit approach. In: Shan, C. (ed.) *APLAS 2013*. LNCS, vol. 8301, pp. 241–256. Springer, Heidelberg (2013)
29. Lax, P.D.: *Functional Analysis*. Pure and Applied Mathematics. Wiley, New York (2002)
30. Kubrusly, C.S.: *The Elements of Operator Theory*, 2nd edn. Birkhäuser, Boston (2011)
31. Greub, W.H.: *Linear Algebra*, vol. 97. Springer, Heidelberg (1967)
32. Goebel, K., Kirk, W.: *Topics in Metric Fixed Point Theory*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, Cambridge (1990)
33. Roman, S.: *Advanced Linear Algebra*. Graduate Texts in Mathematics, vol. 135, 2nd edn. Springer, New York (2005)
34. Kadison, R., Ringrose, J.: *Fundamentals of the Theory of Operator Algebras: Elementary Theory*. AMS (1997). Reprint from Academic Press edition 1983
35. Cousot, P., Cousot, R.: Systematic design of program transformation frameworks by abstract interpretation. In: Proceedings POPL 2002, pp. 178–190 (2002)
36. Cousot, P.: Constructive design of a hierarchy of semantics of a transition system by abstract interpretation. *Theoret. Comput. Sci.* **277**(1–2), 47–103 (2002)
37. Monniaux, D.: Abstract interpretation of probabilistic semantics. In: Palsberg, J. (ed.) *SAS 2000*. LNCS, vol. 1824, pp. 322–339. Springer, Heidelberg (2000)
38. Di Pierro, A., Hankin, C., Wiklicky, H.: Abstract interpretation for worst and average case analysis. In: Reps, T., Sagiv, M., Bauer, J. (eds.) *Wilhelm Festschrift*. LNCS, vol. 4444, pp. 160–174. Springer, Heidelberg (2007)
39. Kocher, P.C.: Timing attacks on implementations of Diffie-Hellman, RSA, DSS, and other systems. In: Koblitz, N. (ed.) *CRYPTO 1996*. LNCS, vol. 1109, pp. 104–113. Springer, Heidelberg (1996)
40. Eaton, J.W., Bateman, D., Hauberg, S.: *GNU Octave - A High-Level Interactive Language For Numerical Computations*, 3rd edn. Cambridge University Press, New York (2007)
41. Crassidis, J.L., Junkins, J.L.: *Optimal Estimation of Dynamic Systems*. Chapman & Hall/CRC Applied Mathematics & Nonlinear Science. CRC Press, Boca Raton (2004)
42. Verhaegen, M., Verdult, V.: *Filtering and System Identification: A Least Squares Approach*. Cambridge University Press, New York (2007)
43. Rao, C.R., Toutenburg, H., Shalabh, Heumann, C.: *Linear Models and Generalizations: Least Squares and Alternatives*. Springer Series in Statistics. Springer, Heidelberg (2008)