Long-Run Cost Analysis by Approximation of Linear Operators over Dioids

David Cachera, Thomas Jensen, Arnaud Jobin, and Pascal Sotin

Irisa, Campus de Beaulieu, 35042 Rennes, France

Abstract. We present a static analysis technique for modeling and approximating the long-run resource usage of programs. The approach is based on a quantitative semantic framework where programs are represented as linear operators over dioids. We provide abstraction techniques for such linear operators which make it feasible to compute safe over-approximations of the long-run cost of a program. A theorem is proved stating that such abstractions yield correct approximations of the program's long-run cost. These approximations are effectively computed as the eigenvalue of the matrix representation of the abstract semantics. The theoretical developments are illustrated on a concrete example taken from the analysis of the cache behaviour of a simple bytecode language.

1 Introduction

This article is concerned with the semantics-based program analysis of quantitative properties pertaining to the use of resources (time, memory, ...). Analysis of such non-functional properties relies on an operational model of program execution where the cost of each computational step is made explicit. We take as starting point a standard small-step operational semantics expressed as a transition relation $\sigma \rightarrow^{q} \sigma'$ between states $\sigma, \sigma' \in \Sigma$ extended with *costs* $q \in Q$ associated to each transition. The set Q of costs is supposed to have two operations for composing costs: a "product" operator that combines the costs along an execution path, and a "sum" operator that combines costs coming from different paths. These operators will give Q a structure of dioid. The sum operator induces a partial order on costs that will serve as a basis for approximating costs. From such a rule-based semantics, there is a straightforward way to obtain a transition matrix, which entries represent the cost of passing from one state of the program to another. This expresses the semantics of a program as a linear operator on $Q(\Sigma)$, the moduloid of vectors of elements of Q indexed over Σ.

In this paper, we are interested in analysing programs with cyclic behaviour (such as reactive systems) in which the asymptotic average cost along cycles, rather than the global cost of the entire execution, is of interest. We define the notion of *long-run cost* for a program which provides an over-approximation of the average cost per transition of long traces. This notion corresponds to the maximum average of costs accumulated along a cycle of the program semantics and is computed from the traces of the successive iterates of the cost matrix. The

quantitative operational semantics operates on state spaces that may be large or even infinite so the computation of quantitative semantic models, like their qualitative counterparts, is usually not tractable. Hence, it is necessary to develop techniques for abstracting this semantics, in order to return an approximation of the program costs that is feasible to compute.

In line with the semantic machinery used to model programs, abstractions are also defined as linear operators from the moduloid over the concrete state space into the moduloid over the abstract one. Given such an abstraction over the semantic domains, we then have to abstract the transition matrix of the program itself into a matrix of reduced size. We give a sufficient condition for an abstraction of the semantics to be correct, *i.e.* to give an over-approximation of the real cost, and show how an abstract semantics that is correct by construction can be derived from the concrete one. The long-run cost of a program is thus safely approximated by an abstract long-run cost, with respect to the order relation induced by the summation operator of the dioid.

The framework proposed here covers a number of different costs related to resource usage (time and memory) of programs. To demonstrate the generality of the framework, our running example considers the less common (compared to time and space) analysis of cache behaviour and the number of cache misses in programs. We illustrate the notions of quantitative semantics, abstraction and long-run cost on a program written in a simple, intermediate bytecode language (inspired by Java Card) onto which we impose a particular cache model.

The paper is structured as follows. Section 2 defines the quantitative semantics as a linear operator over a moduloid. We give the general form of this semantics, and precisely define the notion of cost dioid we use throughout the paper. Section 3 defines the notion of abstraction together with its correctness, and shows how we can derive an abstract semantics that is correct by construction. Section 4 defines the notion of long-run cost, relating it to the asymptotic behaviour of the trace semantics, and shows how a correct abstraction yields an over-approximation of the concrete long-run cost of a program. Section 5 lists related work and Section 6 concludes and discusses future research directions.

2 Linear operator semantics

We give a general framework for expressing quantitative operational semantics. Transitions of these semantics will be equipped with *quantities* (or *costs*) depending on the accessed states. Let P be a program; its semantic domain is the countable set of states Σ. The quantitative operational semantics of P is given as a transition relation, defined by inference rules of the following form: $\sigma \rightarrow^{q} \sigma'$ where σ, σ' are states of Σ, and q is the cost attached to the transition from σ to σ' (q is function of σ and σ'). The set Q of costs and its structure will be made precise in the next subsection. We associate to P the transition system $T = \langle \rightarrow^{\cdot}, I \rangle$, where I is the set of initial states of P. The trace semantics of P is defined as the trace semantics of T.

$$[\![P]\!]_{tr} = [\![T]\!]_{tr} = \{\sigma_0 \rightarrow^{q_0} \ldots \sigma_{n-1} \rightarrow^{q_{n-1}} \sigma_n \mid \sigma_0 \in I, \sigma_i \rightarrow^{q_i} \sigma_{i+1}\}$$

2.1 Cost dioid

The small-step, quantitative operational semantics induces a labelled transition system over Σ with labels in Q and a transition relation $\rightarrow^{\cdot} \subseteq \Sigma \times \Sigma \rightarrow Q$, written $\sigma \rightarrow^{q} \sigma'$. Such a transition states that a direct (one-step) transition from σ to σ' costs q. These unitary transitions can be combined into big-step transitions, using two operators: $\otimes$ for accumulating costs and $\oplus$ to get a maximum of different costs. These operators will form a dioid on Q, as explained below. Costs can be defined in more general ways (for instance, one could use a more general algebra of costs as in [4]) but the present definition covers a number of different costs and has interesting computational properties, since it can be used within a linear operator semantic framework, as presented in the next subsection.

The operator $\otimes$ on Q defines the global cost of a sequence of transitions, $\sigma \rightarrow^{q_1} \ldots \rightarrow^{q_n} \sigma'$ simply as $q = q_1 \otimes \ldots \otimes q_n$. This is written $\sigma \overset{\pi}{\Rightarrow}^{q} \sigma'$ where π is a sequence of states that has σ (resp. σ') as first (resp. last) state.

There may be several ways to reach a state σ' from a state σ, due to the presence of loops and non-determinism in the semantics. Let the set of possible paths be $\Pi_{\sigma,\sigma'} = \{\pi \mid \sigma \overset{\pi}{\Rightarrow}^{q_\pi} \sigma'\}$. The global cost between σ and σ' will be defined, using the operator $\oplus$ on Q, to be $q = \bigoplus_{\pi \in \Pi_{\sigma,\sigma'}} q_\pi$. Formally, the two operators have to fulfill the conditions of a (commutative) dioid.

Definition 1. *A* commutative dioid *is a structure* $(Q, \oplus, \otimes)$ *such that*

1. *Operator $\otimes$ is associative, commutative and has a neutral element e. Quantity e represents a transition that costs nothing.*
2. *Operator $\oplus$ is associative, commutative and has $\bot$ as neutral element. Quantity $\bot$ represents the impossibility of a transition.*
3. *$\otimes$ is distributive over $\oplus$, and $\bot$ is absorbing element for $\otimes$ ($\forall x. x \otimes \bot = \bot \otimes x = \bot$).*
4. *The preorder defined by $\oplus$ ($a \leq b \Leftrightarrow \exists c : a \oplus c = b$) is an order relation (*i.e. *it satisfies $a \leq b$ and $b \leq a \Rightarrow a = b$).*

By nature, a dioid cannot be a ring, since there is an inherent contradiction between the fact that $\oplus$ induces an order relation and the fact that every element has an inverse for $\oplus$. The following lemma is a classical result of dioid theory [17].

Lemma 1. *$\oplus$ and $\otimes$ preserve the order $\leq$, i.e., for all $a, b, c \in Q$ with $a \leq b$, $a \otimes c \leq b \otimes c$ and $a \oplus c \leq b \oplus c$.*

If several paths go from some state σ to a state σ' at the same cost q, we will require that the global cost is also q, *i.e.* we work with idempotent dioids.

Definition 2. *A dioid $(Q, \oplus, \otimes)$ is* idempotent *if $q \oplus q = q$ for all q in Q.*

For instance, $(\overline{\mathbb{R}}, \mathtt{max}, +)$ and $(\overline{\mathbb{R}}, \mathtt{min}, +)$ are idempotent dioids, where $\overline{\mathbb{R}}$ stands for $\mathbb{R} \cup \{-\infty, +\infty\}$. The induced orders are, respectively, the orders $\leq$ and $\geq$ over real numbers, extended to $\overline{\mathbb{R}}$ in the usual way. Note that in an idempotent dioid $a \leq b \Leftrightarrow a \oplus b = b$. Idempotent dioids are also called tropical semirings in the literature. The fact that sets of states may be infinite, together with the use of residuation theory in Section 3 impose that our dioids are complete [7].

Definition 3. *An idempotent dioid is* complete *if it is closed with respect to infinite sums (where we see the operator $\oplus$ as a least upper bound), and the distributivity law holds also for an infinite number of summands.*

A complete dioid is naturally equipped with a top element, that we shall write $\top$, which is the sum of all its elements. Remark that a complete dioid is always a complete lattice, thus equiped with a meet operator $\wedge$ [6]. The notion of long-run cost we will define in Section 4 relies on the computation of an average cost along the transitions of a cycle. This requires the existence of a nth root function.

Definition 4. *A dioid $(Q, \oplus, \otimes)$ is equipped with a* nth root function *if for all q in Q, equation $X^n = q$ has a* unique *solution in Q, denoted by $\sqrt[n]{q}$.*

A sequence containing n transitions, each costing, on the average, $\sqrt[n]{q}$, will thus cost q. Some examples of nth root can be found in Figure 1. To be able to easily deal with the nth root, we make the assumption that the nth power is $\oplus$-lower-semicontinuous ($\oplus$-lsc for short).

Definition 5. *In a complete dioid Q, the nth power is said to be* $\oplus$-lsc *if for all $X \subseteq Q$, $(\bigoplus_{x\in X} x)^n = \bigoplus_{x\in X} x^n$.*

This assumption and its consequences will be very useful for the theorems relating long-run cost and trace semantics in Section 4. Note that this equality remains true for finite X (in that case the nth power is said a $\oplus$-*morphism*).

The following definition summarizes the required conditions for our structure.

Definition 6 (Cost dioid). *A* cost diod *is a* complete *and* idempotent *commutative dioid, equipped with an nth root operation, where the nth power is $\oplus$-lsc.*

Proposition 1. *In a cost dioid Q, we have:*

(i) The nth root is $\oplus$-lsc: $\forall X \subseteq Q, \forall n > 0,\ \sqrt[n]{\bigoplus_{x\in X} x} = \bigoplus_{x\in X} \sqrt[n]{x}$,

(ii) For all $a, b \in Q$ and $n, m > 0$, $\sqrt[n]{a} \oplus \sqrt[m]{b} \geq \sqrt[n+m]{a \otimes b}$.

Property (i) immediately follows from the fact that the nth power is $\oplus$-lsc whereas an intermediate lemma is needed to prove property (ii) [10].

Although the definition of cost dioids may seem rather restrictive, we now show that many classes of dioids found in the literature are indeed cost dioids. We first recall some standard definitions.

Definition 7. *A dioid $(Q, \oplus, \otimes)$ is:*

- selective *if for all a, b in Q, $a \oplus b =$ either a or b.*
- double-idempotent *if both $\oplus$ and $\otimes$ are idempotent.*
- cancellative *if for all a, b, c in Q, $a \otimes b = a \otimes c$ and $a \neq \bot$ imply $b = c$.*

Note that in a double-idempotent dioid, $x^n = x$. Thus, a double-idempotent dioid is naturally equipped with a nth root, which is the identity function.

Proposition 2. *The following dioids are cost dioids.*

(1) Complete and selective commutative dioids with an n*th root operation.*
(2) Complete and double-idempotent commutative dioids.
(3) Complete idempotent commutative dioids satisfying the cancellation condition, and for which for all q *in* Q*, equation* $X^n = q$ *has at least one solution.*

For dioids of kind (1) and (2) we only have to prove that the nth power is $\oplus$-lsc. For dioids of type (3) we also have to prove that if equation $X^n = q$ has a solution, this solution is unique [10].

For instance, $(\overline{\mathbb{R}}, \mathtt{max}, +)$ is a cost dioid that may be used for the definition of the Worst Case Execution Time: when two states can be joined by several sequences of transitions which cost different times, the worst time is taken. To compute the cost of a sequence of transitions, we sum the costs of each transition. Figure 1 lists some examples of cost dioids.

	carrier set	$\oplus$	$\otimes$	$\sqrt[n]{q}$
Double-idempotent	$\mathbb{Q} \cup \{+\infty, -\infty\}$	min	max	q
	$\mathbb{R} \cup \{+\infty, -\infty\}$	max	min	q
	$\mathcal{P}(S)$	$\cap$	$\cup$	q
	$\mathcal{P}(S)$	$\cup$	$\cap$	q
Cancellative	$\mathbb{R}_+^m \cup \{+\infty\}$	min	$+$	$\frac{q}{n}$
Selective	$\mathbb{R}_+ \cup \{+\infty\}$	max	$\times$	$q^{\frac{1}{n}}$
	$\mathbb{Q} \cup \{+\infty, -\infty\}$	max	$+$	$\frac{q}{n}$
	$\mathbb{R} \cup \{+\infty, -\infty\}$	min	$+$	$\frac{q}{n}$

Fig. 1. Some examples of cost dioids

2.2 Semantics as linear operators over dioids

The upshot of using the adequate cost dioid is that the cost computation can be defined in terms of matrix operations in this dioid. The set of one-step transitions can be equivalently represented by a *transition matrix* $M \in \mathcal{M}_{\Sigma\times\Sigma}(Q)$ with

$$M_{\sigma,\sigma'} = \begin{cases} q \text{ if } \sigma \rightarrow^q \sigma' \\ \bot \text{ otherwise} \end{cases}$$

Here, $\mathcal{M}_{\Sigma\times\Sigma}(Q)$ stands for the set of matrices with rows and columns indexed over Σ, and values in Q. This set of matrices is naturally equipped with two operators $\oplus$ and $\otimes$ in the classical way: operator $\oplus$ is extended pointwise, and operator $\otimes$ corresponds to the matrix product (note that the iterate M^n embed the costs for paths of length n). The resulting structure is also an idempotent and complete dioid. The order induced by $\oplus$ corresponds to the pointwise extension of the order over Q: $M \leq M' \Leftrightarrow \forall i, j. M_{i,j} \leq M'_{i,j}$. A transition matrix may also be seen as a linear operator on the moduloid $Q(\Sigma)$, as defined below.

Definition 8. *Let* $(E, \oplus, \otimes)$ *be a commutative dioid. A moduloid over* E *is a set* V *with an internal operation* $\oplus$ *and an external operation* $\odot$ *such that*

1. $(V, \oplus)$ *is a commutative monoid, with* 0 *as neutral element;*
2. *the* $\odot$ *operator ranges from* $E \times V$ *to* V*, and verifies*
   (a) $\forall \lambda \in E, \forall (x, y) \in V^2, \lambda \odot (x \oplus y) = (\lambda \odot x) \oplus (\lambda \odot y)$,
   (b) $\forall (\lambda, \mu) \in E^2, \forall x \in V, (\lambda \oplus \mu) \odot x = (\lambda \odot x) \oplus (\mu \odot x)$,
   (c) $\forall (\lambda, \mu) \in E^2, \forall x \in V, \lambda \odot (\mu \odot x) = (\lambda \otimes \mu) \odot x$,
   (d) $\forall x \in V, e \odot x = x$ *and* $\perp \odot x = 0$,
   (e) $\forall \lambda \in E, \lambda \odot 0 = 0$.

If E is an idempotent dioid, then any moduloid V over E is also an idempotent dioid, equipped with a canonical order defined from the $\oplus$ operation. As for vector spaces, if n is a given integer, E^n, set of vectors with n components in E, is a moduloid. More generally, a vector $u \in E(\Sigma)$, with Σ finite, $|\Sigma| = n$ can be seen as a function $\delta_u : [1, n] \to E$. Since Q is complete, we can generalize to the infinite (countable) case: δ_u becomes a mapping from $\mathbb{N}$ to E, and the same technique applies for matrices. The matrix-vector product is defined by: $(Mu)_i = \bigoplus_{j=1}^{+\infty} \delta_M(i, j) \otimes \delta_u(j)$. In this paper, we will keep the matrix notation for the sake of simplicity, even for an infinite set of indices.

2.3 Running example: quantitative semantics

We illustrate the notions of cost and quantitative semantics on a simple bytecode language, inspired by the Java Card language. Figure 2 shows part of the factorial program written in this language. The quantity we are interested in is the number of cache misses related to read accesses (read miss behaviour). In order to describe the read miss behaviour of programs, we extend the semantics of a simple bytecode language [19,16] with a cache model and with quantities expressing the number of read misses.

Source	Bytecode
x=1;	`1: push 1`
	`2: store x`
for (i=2;...	`3: push 2`
	`4: store i`
...i<=n;...	`5: load i`
	`6: load n`
	`7: if ≤goto 14`
*x=x*i;*	`8: load x`
	`9: load i`
	`10: numop mul`
	`11: store x`
...i++)	`12: inc i`
	`13: goto 5`
return x;	`14: load x`
	`15: return`

Fig. 2. Factorial program

The cost dioid considered here is $(\overline{\mathbb{R}}, \mathtt{max}, +)$. A state contains a heap, a call stack of frames, and within each frame an instruction pointer for the current method, an array of local variables and an operand stack. In addition to these standard elements, a state contains a set of *logical addresses*, representing which values are present in the cache at this point of the execution. This set is managed similarly to the cache. For example, the maximum size of this set will correspond to the size of the physical cache, and the replacement policy will model the one it provides (*e.g.* LRU, FIFO). The cache description is hidden in a function $C' = \mathit{update}(C, [\mathit{access}])$ where C and C' denote the cache before and after a transition, respectively, and where $[\mathit{access}]$ is a list of memory accesses. Due to the lack of space, full descriptions of possible *update* functions are not given.

Memory is accessed with two operators (a read or write access) which take two parameters, specifying what volume of data is to be accessed, and where these data are stored. For example, $read_\tau(\texttt{heap}.\mathit{3}.x)$ means that data of type τ is read at the address $\texttt{heap}.\mathit{3}.x$, *i.e.* field x of the third object in the heap. In the same way, $\texttt{stack}.\mathit{frameId}.n$ points to the n-nth element in the operand stack of a given frame, and $\texttt{local}.\mathit{frameId}.\mathit{local}$ points to a local variable in a certain frame. We give an example of a semantic rule: the `load` instruction, which loads a typed local variable, indexed by i, on the top of the operand stack. The first two hypotheses of the rule correspond to the standard semantics. The third and fourth hypotheses define how the cache evolves when executing a `load`. The fifth hypothesis computes the cost. Some other rule examples can be found in [10].

$$\frac{\begin{array}{c} \mathit{InstrAt}(m, ip) = \texttt{load}\ \tau\ \ i \wedge L[i] = d \\ S' = d :: S \wedge \mathit{size}(S) = t \\ \mathit{access} = \texttt{[}\mathit{read}_\tau(\texttt{local}.f.i); \mathit{write}_\tau(\texttt{stack}.f.t+1)\texttt{]} \\ C' = \mathit{update}(C, \mathit{access}) \\ q = \mathit{nbRmiss}(C, \mathit{access}) \end{array}}{\langle H, \ll f, m, ip, L, S \gg :: fr, C\rangle \to^q \langle H, \ll f, m, ip+1, L, S' \gg :: fr, C'\rangle}$$

The number of read misses depends on the current state of the cache and the way it is accessed. This is defined precisely by the function $\mathit{nbRmiss}(C, \mathit{access})$ that computes the number of read misses generated by the list of memory accesses *access* if the cache at the beginning of the instruction is C. Here is the pseudocode of function *nbRmiss*.

$$\mathit{nbRmiss}(c,\texttt{[]}) = 0$$

$$\mathit{nbRmiss}(c,\texttt{[}a\,|\,r\texttt{]}) = \mathit{nbRmiss}(\mathit{update}(c,\texttt{[}a\texttt{]}),r) + \begin{cases} 1 & \mathit{if}\ \ a = \mathit{read}\ m\ \ \mathit{and}\ \ m \notin c \\ 0 & \mathit{otherwise} \end{cases}$$

3 Abstraction

The transition matrix representing a program is in general of infinite dimension, so neither transitive closure nor traces can be computed in finite time. To overcome this problem, we define an abstract matrix that can be used to approximate the computations of the original matrix. For example, if we compute the minimum memory needed to run a program, a correct approximation of this quantity must be greater than the effective minimum. In this section, we give a sufficient condition for this approximation to be correct with respect to the ordering induced by the dioid. To prove the correctness of an abstraction, we restate the classical abstract interpretation theory [12] in terms of linear operators over moduloids.

3.1 Galois connections and pseudo-inverses

We first briefly recall the definition of Galois connections that are used in the classical abstract interpretation theory.

Definition 9. *Let $(C, \leq_C)$ and $(D, \leq_D)$ be two partially ordered sets. Two mappings $\alpha : C \mapsto D$ (called abstraction function) and $\gamma : D \mapsto C$ (called concretization function) form a Galois connection (C, α, γ, D) iff:*

- $\forall c \in C, \forall d \in D, c \leq_C \gamma(d) \iff \alpha(c) \leq_D d$, *or equivalently*
- α *and* γ *are monotonic and* $\alpha \circ \gamma \leq Id_D$ *and* $Id_C \leq \gamma \circ \alpha$

In our setting, the partial orders will be the orders induced by the $\oplus$ operators over vectors in a moduloid. The question that naturally arises is that of the existence of a concretization function, given an abstraction α. In [15], Di Pierro and Wiklicky describe the framework of Probabilistic Abstract Interpretation, where the abstraction function is a linear operator over the semiring of probabilities. They obtain a concretization function through the Moore-Penrose pseudo-inverse. As we will not be able to define an exact inverse in the general case, nor to apply the Moore-Penrose pseudo-inverse since we do not work in a field, we will use the theory of *residuation* to get a kind of inverse for α.

Proposition 3. *Let E and F be two complete dioids, f a monotone mapping from E to F. We call subsolution of equation $f(x) = b$ an element y such that $f(y) \leq b$. The following properties are equivalent.*

1. *For all $b \in E$, there exists a greatest subsolution to the equation $f(x) = b$.*
2. *$f(\bot_E) = \bot_F$, and f is $\oplus$-lsc.*
3. *There exists a monotone mapping from F into E which is upper*[1] *semi-continuous such that $f \circ f^\dagger \leq Id_F$ and $Id_E \leq f^\dagger \circ f$.*

Consequently, $f^\dagger$ is unique. When f satisfies these properties, it is said to be residuated, and $f^\dagger$ is called its residual.

3.2 Abstraction over cost dioids

We now show how the notions of abstraction and concretization can be recast in our setting. In the following, Σ will denote a set of *concrete* states and $\Sigma^\sharp$ a set of *abstract* states. An abstraction function maps concrete states in Σ to their abstraction in $\Sigma^\sharp$. Given an abstraction function α, we can lift it to a linear abstraction operator $\alpha^\uparrow \in \mathcal{M}_{\Sigma^\sharp \times \Sigma}(Q)$ by setting

$$\alpha^\uparrow_{\sigma^\sharp,\sigma} = \begin{cases} e \text{ if } \alpha(\sigma) = \sigma^\sharp \\ \bot \text{ otherwise} \end{cases}$$

In what follows, $\alpha^\uparrow$ will be denoted by α when no confusion can arise and $\leq$ will stand for the order defined on $\mathcal{M}_{\Sigma^\sharp \times \Sigma^\sharp}(Q)$ in Section 2.2.

Recall that the set of linear operators over a complete idempotent moduloid is itself a complete idempotent dioid. As the abstraction function is linear, it trivially fulfills requirements 2 of Proposition 3 and we get the following result.

Theorem 1. *Let Σ and $\Sigma^\sharp$ be the domains of concrete and abstract states, α a mapping from Σ to $\Sigma^\sharp$, and $\alpha^\uparrow \in \mathcal{M}_{\Sigma^\sharp \times \Sigma}(Q)$ the linear mapping obtained by lifting α. There exists a unique monotonic $\alpha^\dagger$ such that $\alpha^\uparrow \circ \alpha^\dagger \leq Id_{\Sigma^\sharp}$ and $Id_\Sigma \leq \alpha^\dagger \circ \alpha^\uparrow$.*

[1] Upper semi-continuity is the analog of lower semi-continuity for the $\wedge$ operator.

3.3 Induced abstract semantics

Let T be a transition system in the concrete domain Σ, over the cost dioid $(Q, \oplus, \otimes)$. We now want to define an abstract transition system over the abstract domain $\Sigma^\sharp$ that is "compatible" with T, both from the point of view of its traces and from the costs it will lead to compute. The following definition of a correct abstraction will ensure that the long-run cost of a program, as defined in the next section, will be correctly over-approximated during the abstraction process.

Definition 10 (Correct abstraction). *Let $T = \langle M, I \rangle$ a transition system over the concrete domain, with $M \in \mathcal{M}_{\Sigma\times\Sigma}(Q)$ and $I \subseteq \Sigma$. Let $T^\sharp = \langle M^\sharp, I^\sharp \rangle$ be a transition system over the abstract domain, with $M^\sharp \in \mathcal{M}_{\Sigma^\sharp\times\Sigma^\sharp}(Q)$ and $I^\sharp \subseteq \Sigma^\sharp$. Let α be an abstraction from Σ to $\Sigma^\sharp$. The triple $(T, T^\sharp, \alpha)$ is a correct abstraction from Σ to $\Sigma^\sharp$ if $\alpha^\uparrow \circ M \leq M^\sharp \circ \alpha^\uparrow$ and $\{\alpha(\sigma) \mid \sigma \in I\} \subseteq I^\sharp$.*

The classical framework of abstract interpretation gives a way to define a best correct abstraction for a given concrete semantic operator. In the same way, given an abstraction α and a concrete semantics linear operator, we can define an abstract semantics operator that is correct by construction, as expressed by the following proposition.

Proposition 4. *Let α be an abstraction from Σ to $\Sigma^\sharp$, and $T = \langle M, I \rangle$ be a transition system with $M \in \mathcal{M}_{\Sigma\times\Sigma}(Q)$ a linear operator over the concrete moduloid and I the subset of initial states. We set $T^\sharp = \langle M^\sharp, I^\sharp \rangle$ with*

$$M^\sharp = \alpha^\uparrow \circ M \circ \alpha^\dagger \quad \text{and} \quad I^\sharp = \{\alpha(\sigma) \mid \sigma \in I\}$$

Then $(T, T^\sharp, \alpha)$ is a correct abstraction from Σ to $\Sigma^\sharp$. Moreover, given T and α, $T^\sharp$ provides the best possible abstraction in the sense that if $(T, \langle M', I' \rangle, \alpha)$ is another correct abstraction, then $M^\sharp \leq M'$ and $I^\sharp \subseteq I'$.

Proof. The proof follows from the facts that $Id \leq \alpha^\dagger \circ \alpha$ and $\alpha \circ \alpha^\dagger \leq Id$.

The above definitions and properties deal with the matrix view of the semantics, but what can be said about traces? The following proposition states that for each program trace, there exists an "abstract" trace of same length which costs are given by the induced abstract matrix. This property will be useful for proving the correctness of abstractions in Section 4.

Proposition 5. *Let consider the transition system $T = \langle q, I \rangle$ with $I \subseteq \Sigma$ its set of initial states and $q : \Sigma \times \Sigma \to Q$ its quantitative transition system in the cost dioid Q. Let α be an abstraction function from Σ to $\Sigma^\sharp$. Let $T^\sharp = \langle q^\sharp, I^\sharp \rangle$ an abstract transition system defined by:*

- $I^\sharp = \{\alpha(\sigma) \mid \sigma \in I\}$
- $\alpha^{-1} : \Sigma^\sharp \to \mathcal{P}(\Sigma)$ *with* $\alpha^{-1}(\sigma^\sharp) = \{\sigma \mid \alpha(\sigma) = \sigma^\sharp\}$
- $q^*(\Sigma_1, \Sigma_2) = \bigoplus_{(\sigma_1,\sigma_2)\in\Sigma_1\times\Sigma_2} q(\sigma_1, \sigma_2)$
- $q^\sharp(\sigma_1^\sharp, \sigma_2^\sharp) = q^*(\alpha^{-1}(\sigma_1^\sharp), \alpha^{-1}(\sigma_2^\sharp))$

then forall $t = \sigma_0 \to^{q_0} \ldots \sigma_n \in [\![T]\!]_{tr}, |t| = n$, there exists $t^\sharp = \sigma_0^\sharp \to^{q_0'} \ldots \sigma_n^\sharp \in [\![T^\sharp]\!]_{tr}, |t| = n$ such that $q_i \leq q_i' \quad \forall i \in [0, n-1]$ and $\sigma_i^\sharp = \alpha(\sigma_i) \quad \forall i \in [0, n]$. In addition, $M^\sharp = \alpha^\uparrow \circ M \circ \alpha^\dagger$ is the transition matrix for $q^\sharp$.

3.4 Running example: abstraction

In 2.3, we introduced a quantitative semantics describing the number of cache misses in read access. M is the matrix describing this quantitative semantics for the factorial program (see Figure 2 for the code). The exact computation of the semantics would be too costly, even if we work with bounded numerical domains. In this subsection, we are using the abstractions techniques in order to compute an abstract semantics $M^\sharp$ from the matrix M.

We abstract a concrete state by the instruction pointer and the k last data accessed. Within this abstract domain, the loss of information lies in three points:

- Values (*i.e.* locals, stack and heap) are forgotten. This prevents us from determining the value of branching condition.
- The cache size is reduced to k elements. When k grows, precision increase, and so do the cost of the analysis.
- The method call stack is forgotten. We turn the analysis into an intra-procedural one, not for efficiency but for clearer notation, as our factorial function involves only one non-recursive function.

We write the abstract state as $(ip, [v_1, \dots, v_k])$ where ip is the instruction pointer and $[v_1, \dots, v_k]$ is a list of logical addresses of the last data accessed, v_k being the most recent. $\mathtt{s}.0$ refers to the bottom element of the local stack, $\mathtt{l}.x$ refers to the local variable called x in the source code. We model the maximum number of read misses using the dioid $Q = (\overline{\mathbb{R}}, \mathtt{max}, +)$.

$$M^\sharp = \left(\begin{array}{c|ccccc} & \dots & 9, [\mathtt{l}.x, \mathtt{s}.0] & \dots & 9, [\mathtt{l}.i, \mathtt{l}.x, \mathtt{s}.0] & \dots \\ \hline \vdots & & & & & \\ 8, [] & & 1 & & \bot & \\ 8, [\mathtt{l}.x] & & 0 & & \bot & \\ 8, [\mathtt{s}.0] & & 1 & & \bot & \\ 8, [\mathtt{s}.0, \mathtt{l}.x] & & 0 & & \bot & \\ 8, [\mathtt{l}.x, \mathtt{s}.0] & & 0 & & \bot & \\ 8, [\mathtt{l}.i, \mathtt{l}.x] & & \bot & & 0 & \\ 8, [\mathtt{l}.i] & & \bot & & 1 & \\ \vdots & & & & & \end{array}\right)$$

Fig. 3. Transition matrix

We construct the abstract matrix associated to our abstract system. Its size is bounded in terms of the cardinality of $\mathcal{I}$, the set of all instruction pointers appearing in program P, and the number of up-to-k-combinations of the different logical data used in this function (which form a finite set $\mathcal{L}$). A value $q^\sharp$ of this matrix, standing at row $a^\sharp$ and column $b^\sharp$ ($a^\sharp$ and $b^\sharp$ are two abstract states), is computed in this way: let A and B be the set of concrete states abstracted by $a^\sharp$ and $b^\sharp$. Then $q^\sharp = \bigoplus\{q \mid a \to^q b, a \in A, b \in B\}$. For example

- $\sigma = (8, [\mathtt{l}.x]) \to^0 (9, [\mathtt{l}.x, \mathtt{s}.0]) = \sigma'$,
  Whatever the concrete state and its precise values, if it is abstracted by σ, then it can turn into a state abstracted by σ' for a cost of 0 read miss.
- $\sigma = (8, [\mathtt{s}.0]) \to^1 (9, [\mathtt{l}.x, \mathtt{s}.0]) = \sigma'$,
  In the same way, all states abtracted by σ can generate up to one read miss on their next instruction, turning into states abstracted by σ'.

Recall that the `load` x instruction at line 8 accesses these memory locations: $[read_\tau(\texttt{local}.f.i); write_\tau(\texttt{stack}.f.t[+1])]$ with i the local variable, t the current stack height and f the current frame.

A $\perp$-transition denotes an incompatibility between the two abstract states, either in its control flow or its the cache evolution. Most of the matrix will be filled by $\perp$. This kind of matrix is called sparse matrix, and permits the use of particularly small representations together with efficient algorithms. Figure 3 gives a submatrix of the abstract matrix. $M^\sharp \in \mathcal{M}_{(\mathcal{I}\times\{\emptyset\cup\mathcal{L}\cup\mathcal{L}^2\})^2}(Q)$.

4 Long-run cost

So far, we have seen that all single-transition costs can be summarized in a transition matrix. We now use this matrix and the mathematical results of dioid algebra to define a notion of long run cost for a whole program. In [20] we proposed a notion of *global cost* of a program, representing its cost from initial to final states. It correctly deals with programs which are meant to terminate, but in some cases this global cost turns out to be $\top$, in particular when it is evaluated on a coarse abstraction of the initial system. Getting $\top$ as a result for the global cost is rather unsatisfactory as it does not tell anything about the concrete cost. For this case and for the case of programs which are not meant to terminate (as reactive systems), we propose the notion of *long-run cost*, that represents a maximal average cost over cycles of transitions. This terminology is taken from [2,9], in the context of probabilistic processes modelled by Markov decision processes. Behaviour patterns of interest (described by labelled graphs) are associated to real numbers representing the success or the duration of the pattern, and extensions of branching time temporal logics are proposed in order to measure their long-run average outcome.

The average cost of a finite path is defined as the arithmetical mean (w.r.t. the $\otimes$ operator) of the costs labelling its transitions. In other words, it is the nth root of the global cost of the path, where n is its length. We write $\tilde{q}(\pi) = \sqrt[|\pi|]{q(\pi)}$ for the average cost of path π, where $q(\pi)$ is the global cost of π, and $|\pi|$ its length. The "maximum" average cost of all cycles in the graph will be the quantity we are interested in: this quantity will be called *long-run cost*. The following example illustrates these notions on a simple graph.

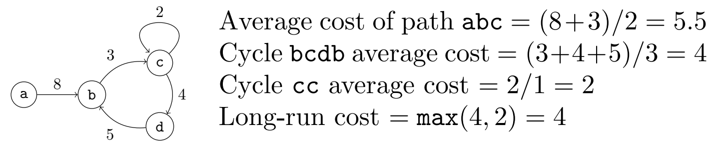

By the properties of the dioids we consider, matrix M^k sums up the transition costs of all paths of length k. The diagonal of this matrix thus contains the costs of all cycles of length k. If we add up all the elements on this diagonal, we get the trace of the matrix. This observation gives rise to the following definition.

Definition 11. *Let P be a program having $T = \langle M, I \rangle$ for transition system. Let R be M restricted to the set of states, Σ, reachable from I. The long-run cost of program P is defined as the long-run cost of T*

$$\rho(P) = \rho(T) = \bigoplus_{k=1}^{|\Sigma|} \sqrt[k]{tr\ R^k} \quad \textit{where} \quad tr\ R = \bigoplus_{1}^{|\Sigma|} R_{i,i}.$$

Note that this definition is valid even for an infinite number of states, since we work with complete dioids. As an example, if we work in the dioid $(\mathit{Time}, \mathtt{max}, +)$, where *Time* is isomorphic to $\overline{\mathbb{R}}$, $\rho(P)$ is the maximal average of time spent per instruction, where the average is computed on any cycle by dividing the total time spent in the cycle by the number of instructions in this cycle. In the case of a finite set of states, the long-run cost is computable, and we note in the passing that its definition coincides with the definition of the maximum of eigenvalues of the matrix, in the case of an irreducible matrix in an idempotent semiring [11].

4.1 Semantics of the long run cost

The following proposition establishes in a more formal manner the link between this definition of long-run cost and the cycles of the semantics.

Proposition 6. *Let Γ be the set of cycles in T. Then $\rho(T) = \bigoplus_{c \in \Gamma} \tilde{q}(c)$.*

The idea of the proof is to show that the cycles of length less than $|\Sigma|$ are enough to know average costs, and that a partition of these cycles is related with the different iterates of the matrix appearing in Definition 11. The proof becomes straightforward in the case of an infinite set of states.

As we aim at giving a characterisation of the asymptotic behaviour of a program, an alternative definition for long-run cost could have been:

$$lrc(T) = \limsup_{n \to \infty} \bigoplus_{\substack{t \in [\![T]\!]_{tr} \\ |t| = n}} \tilde{q}(t)$$

Instead of defining the long-run cost w.r.t. the cycles, this definition considers arbitrarily long traces. Unlike $\rho(P)$, $lrc(P)$ is not suitable for computation, even if the set of states is finite. We will see in Subsection 4.3 that those two notions coincide in a restricted class of cost dioids and when the set of states is finite.

4.2 Ensuring correctness

The question that naturally arises is to know if the notion of long-run cost is preserved by abstraction. The following theorem states that a correct abstraction gives an over-approximation of the concrete long-run cost.

Theorem 2. *If $(T, T^\sharp, \alpha)$ is a correct abstraction, then $\rho(T) \leq_Q \rho(T^\sharp)$.*

The proof of theorem relies on the fact that the correctness is preserved when the concrete and abstract matrices are iterated simultaneously [10].

Recall that Proposition 5 states that for any concrete trace, there exists an abstract trace which cost is over the concrete one. It follows that the alternative definition of long-run cost given in Section 4.1 is also preserved by abstraction:

Proposition 7. *If $(T, T^\sharp, \alpha)$ is a correct abstraction, then*

$$\limsup_{n\to\infty} \bigoplus_{\substack{t\in[\![T]\!]_{tr}\\ |t|=n}} \tilde{q}(t) \leq \limsup_{n\to\infty} \bigoplus_{\substack{t^\sharp\in[\![T^\sharp]\!]_{tr}\\ \left|t^\sharp\right|=n}} \tilde{q}^\sharp(t^\sharp)$$

4.3 Traces meet cycles

We now show that, if Σ is finite, and for dioids where the carrier set is $\overline{\mathbb{R}}$ and operator $\otimes$ is the arithmetical $+$ (so that the nth root operator corresponds to division by n), the notion of long-run cost defined w.r.t. accessible cycles coincides with the notion of long-run cost defined as the limit of the maximum average cost of traces which length tends to infinity. To establish this result, we have to show that the cost of a prefix of a trace becomes negligible when this trace becomes arbitrarily long. We thus impose the following hypothesis:

Hypothesis 1 *All transitions δ which are not in a cycle verify $q(\delta) \neq +\infty$.*

Hypothesis 1 excludes certain pathological matrices with atomic operations that have infinite costs. If a cycle contains a $+\infty$ transition, the ρ value indicates it.

Theorem 3. *Let $T = \langle M, I\rangle$ be a transition system with $M \in \Sigma \times \Sigma \to Q$. If Σ is finite and Q is a cost dioid where the carrier set is $\overline{\mathbb{R}}$ and operation $\otimes$ is the arithmetical $+$, then with Hypothesis 1, we have*

$$\rho(T) = \lim_{n\to\infty} \bigoplus_{\substack{t\in[\![T]\!]_{tr}\\ |t|=n}} \tilde{q}(t)$$

This theorem establishes a link between the semantics and a computable definition of the long-run cost. The key points of the proof are to ensure that this limit exist, and to show that a small part of a trace can be neglected for very long traces. This is proved by bounding $\bigoplus \tilde{q}(t)$ [10].

4.4 Running example: long-run cost

To illustrate the use of the long-run cost (ρ), we will consider a cache which can contain 4 integers. Such a small size could seem weird to the reader and unrealistic for a cache size, but the term cache can be interpreted here as some kind of registers. The semantics of the factorial program is abstracted as described in Section 3.4, with $k = 4$. Using Definition 11, we compute $\rho(M^\sharp) = 2/9$, meaning that in an execution long enough, we observe on average 2 cache misses each 9

instructions. A prototype is under developement for the standard Java bytecode, that already handles this kind of example. It implements the analysis technique presented in this paper with the help of an existing Scilab library for max-plus algebra, that offers an efficient ways to compute the long-run cost (sparse matrices and Howard's algorithm for eigenvalues [11], handle 20000 accessible states).

If we now consider a FIFO replacement policy for the same 4 integer registers, we obtain a different long-run cost. The FIFO policy implementation is cheaper in electronic components than the LRU one, but the analysis of the factorial function says that $\rho = 4/9$, *i.e.*, that we now have on average 4 cache misses for 9 instruction executed. Such a slowdown is coherent with the observations that small cache memory requires more advanced cache policies.

5 Related work

The present work is inspired by the quantitative abstract interpretation framework developed by Di Pierro and Wiklicky [15]. We have followed their approach in modeling programs as linear operators over a vector space, with the notable technical difference that their operators act over a semiring of probabilities whereas ours work with idempotent dioids. Working with idempotent dioids means that we have been able to exploit known results from Discrete Event Systems theory which makes intensive use of such structures. Another difference with respect to [15] lies in the kind of program being analyzed: we have been considering an intermediate bytecode language rather than declarative languages (probablistic concurrent constraint programming and the lambda calculus [14]).

In Di Pierro and Wiklicky's work, the relation with abstract interpretation is justified by the use of the pseudo-inverse of a linear operator, similar to a Galois connection mechanism, enforcing the soundness of abstractions. Our approach can be seen as intermediate between their and classical abstract interpretation: on one hand, we use residuation theory in order to get a pseudo-inverse for linear abstraction functions; on the other hand, we benefit from the partially ordered structure of dioids to give guarantees of soundness under the assumption $\alpha \circ M \leq_D M^\sharp \circ \alpha$, which is a classical requirement in abstract interpretation.

Several other works make use of idempotent semiring for describing quantitative aspects of computations, namely under the form of constraint semirings [8]. Recently, these have been used in the field of Quality of Services [13], in particular with systems modelled by graph rewriting mechanisms [18]. In all these approaches, the $\oplus$ and $\otimes$ operators of the constraint semiring are used for combining constraints. In [5], Aziz makes use of semirings in a mobile process calculus derived from the π-calculus, in order to model the cost of communicating actions. He also defines a static analysis framework, by abstracting "concrete" semirings into abstract semirings of reduced cardinality, and defining abstract semiring operators accordingly. For instance, the $(\mathbb{R}_+ \cup \{+\infty\}, \min, +)$ semiring can be abstracted by a $(\{\mathit{low}, \mathit{medium}, \mathit{high}\}, \min, \max)$ one. Even if dealing with dioids as we do, none of these approaches make use of a notion of long-run cost to express an average quantitative behaviour of a system.

In the specific context of Java bytecode, Albert and al. [1] defined a cost analysis based on the generation of cost relations and recurrence equations. Approximation of costs is done in two steps: first, a classical abstract interpretation is used to approximate size relations between variables. Secondly, combining size relations then gives recurrence equations whose solutions are approximated by using intervals when no closed form solution has been found. This gives interesting results for a class of simple programs, in particular when arithmetic operations are restricted to linear ones.

Our running example of estimating cache usage is meant for illustrative purposes and is based on a rather abstract view of cache analysis, compared *e.g.* to the detailed modeling and cache abstraction of Wilhelm and al. [3] who propose in the AbsInt tool a cache behaviour prediction by abstract interpretation. Three points of their work could be almost directly used in our framework: the various models of cache (*e.g.* direct-mapped, A-way) to implement our update function, their abstract domain, in order to design our quantitative abstractions, and their observations about caches and writing, in order to develop an accurate model. Their approach however is not directed toward long-run cost computation.

6 Conclusion

We have shown how to abstract the long-run cost of programs whose operational semantics is defined as transition systems labelled by costs taken from a particular kind of dioids. In such cases, we have shown that the semantics is a linear operator over the moduloid associated to this dioid. We have used a well-known characterization of the asymptotic behaviour of a discrete event system to define the notion of long-run cost of such a semantics, and proposed a novel way of analyzing the long-run behaviour of the program. We have characterized this long-run cost as being a maximal average cost per transition on very long traces of the semantics. Computing the exact long-run cost of a program is in general too expensive, so we have extended the linear operator framework with a notion of abstraction of the semantics which is also expressed as a linear operator. A correctness relation between concrete and abstract semantic matrices ensures that the cost computed from the abstract semantics is an over-approximation of the concrete one. The notions of dioids, quantitative semantics, abstraction and long-run cost have been illustrated all along the paper through a cache miss analysis on a program written in a simple bytecode language.

Future work The examples in the paper have been computed both by hand (for the abstraction part) and by a prototype analyzer for the computation of long-run costs themselves. Future work includes improvement of the prototype and developement of a framework for validating experimental results.

An interesting avenue for further work would be to relax the correctness criterion so that the abstract estimate is “close” to (but not necessarily greater than) the exact quantity. For certain quantitative measures, a notion of “closeness” might be of interest, as opposed to the qualitative case where static analyses must err on the safe side.

References

1. E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost Analysis of Java Bytecode. In *ESOP'07*, volume 4421 of *Lecture Notes in Computer Science*, pages 157–172. Springer, 2007.
2. L. D. Alfaro. How to Specify and Verify the Long-Run Average Behavior of Probabilistic Systems. In *13th Symposium on Logic in Computer Science (LICS'98)*, pages 174–183. IEEE Computer Society Press, 1998.
3. M. Alt, C. Ferdinand, F. Martin, and R. Wilhelm. Cache Behavior Prediction by Abstract Interpretation. In *Static Analysis Symposium (SAS'96)*, volume 1145 of *Lecture Notes in Computer Science*, pages 51–66, September 1996.
4. D. Aspinall, L. Beringer, M. Hofmann, H.-W. Loidl, and A. Momigliano. A program logic for resources. *Theor. Comput. Sci.*, 389(3):411–445, 2007.
5. B. Aziz. A Semiring-based Quantitative Analysis of Mobile Systems. *Electronic Notes in Theoretical Computer Science*, 157(1):3–21, 2006.
6. F. Baccelli, G. Cohen, G. J. Olsder, and J.-P. Quadrat. *Synchronization and Linearity*. Wiley, 1992.
7. S. Bistarelli and F. Gadducci. Enhancing Constraints Manipulation in Semiring-Based Formalisms. In *European Conf. on Artificial Intelligence*, pages 63–67, 2006.
8. S. Bistarelli, U. Montanari, and F. Rossi. Semiring-Based Constraint Satisfaction and Optimization. *Journal of the ACM*, 44(2):201–236, 1997.
9. T. Brazdil, J. Esparza, and A. Kucera. Analysis and Prediction of the Long-Run Behavior of Probabilistic Sequential Programs with Recursion. In *FOCS '05: Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 521–530, Washington, DC, USA, 2005. IEEE Computer Society.
10. D. Cachera, T. Jensen, A. Jobin, and P. Sotin. Long-run cost analysis by approximation of linear operators over dioids. Research Report 6338, INRIA, 2007.
11. J. Cochet-Terrasson, G. Cohen, S. Gaubert, M. Mc Gettrick, and J.-P. Quadrat. Numerical Computation of Spectral Elements in Max-Plus Algebra. In *Proceedings of the IFAC Conference on System Structure and Control*, 1998.
12. P. Cousot and R. Cousot. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *4th Symposium on Principles of Programming Languages (POPL'77)*, 1977.
13. R. De Nicola, G. Ferrari, U. Montanari, R. Pugliese, and E. Tuosto. A Basic Calculus for Modelling Service Level Agreements. In *Inter. Conf. on Coordination Models and Languages*, volume 3454 of *LNCS*, April 2005.
14. A. Di Pierro, C. Hankin, and H. Wiklicky. Probabilistic λ-calculus and Quantitative Program Analysis. *J. Logic and Computation*, 15(2):159–179, 2005.
15. A. Di Pierro and H. Wiklicky. Concurrent Constraint Programming: Towards Probabilistic Abstract Interpretation. In *PPDP*, 2000.
16. S. N. Freund and J. C. Mitchell. A Formal Framework for the Java Bytecode Language and Verifier. *ACM SIGPLAN Notices*, 34(10):147–166, 1999.
17. M. Gondran and M. Minoux. *Graphes, dioïdes et semi-anneaux*. Tec & Doc, 2001.
18. D. Hirsch and E. Tuosto. SHReQ: Coordinating Application Level QoS. In *SEFM '05: Proceedings of the Third IEEE International Conference on Software Engineering and Formal Methods*, pages 425–434, Washington, DC, USA, 2005.
19. I. Siveroni. Operational Semantics of the Java Card Virtual Machine. *J. Logic and Automated Reasoning*, 2004.
20. P. Sotin, D. Cachera, and T. Jensen. Quantitative Static Analysis over Semirings: Analysing Cache Behaviour for Java Card. In A. Di Pierro and H. Wiklicky, editors, *QAPL06, Quantitative Aspects of Programming Languages*, 2006.