

 Open access • Journal Article • DOI:10.1145/1555746.1555753

Resource control graphs — Source link

Jean-Yves Moyen

Institutions: University of Paris

Published on: 14 Aug 2009 - ACM Transactions on Computational Logic (ACM)

Topics: Program analysis, Termination analysis and Representation (mathematics)

Related papers:

- [Certifying Polynomial Time and Linear/Polynomial Space for Imperative Programs](#)
- [A new recursion-theoretic characterization of the polytime functions](#)
- [Termination and Non-termination in Logic Programming](#)
- [Quasi-terminating logic programs for ensuring the termination of partial evaluation](#)
- [Proving termination with Multiset Orderings](#)

Share this paper:    

View more about this paper here: <https://typeset.io/papers/resource-control-graphs-2iw1tqe19d>



HAL
open science

Resource Control Graphs

Jean-Yves Moyen

► **To cite this version:**

Jean-Yves Moyen. Resource Control Graphs. ACM Transactions on Computational Logic, Association for Computing Machinery, 2009, 10 (4), 10.1145/1555746.1555753 . hal-00107145v3

HAL Id: hal-00107145

<https://hal.archives-ouvertes.fr/hal-00107145v3>

Submitted on 3 Sep 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Resource Control Graphs

JEAN-YVES MOYEN

University of Paris 13

Resource Control Graphs can be seen as an abstract representation of programs. Each state of the program is abstracted as its size, and each instruction is abstracted as the effects it has on the size whenever it is executed. The Control Flow Graph of the programs gives indications on how the instructions might be combined during an execution.

Termination proofs usually work by finding a decrease in some well-founded order. Here, the sizes of states are ordered and such kind of decrease is also found. This allows to build termination proofs similar to the ones in Size Change Termination.

But the size of states can also be used to represent the space used by the program at each point. This leads to an alternate characterisation of the Non Size Increasing programs, that is the ones that can compute without allocating new memory.

This new tool is able to encompass several existing analyses and similarities with other studies hint that even more might be expressible in this framework thus giving hopes for a generic tool for studying programs.

Categories and Subject Descriptors: D.2.4 [Software engineering]: Software/Program Verification; F.2.2 [Analysis of algorithms and problem complexity]: Nonnumerical Algorithms and Problems—*Computations on discrete structures*; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and reasoning about Programs; G.2.2 [Discrete Mathematics]: Graph Theory

General Terms: Algorithms, Theory, Verification

Additional Key Words and Phrases: Abstraction, implicit complexity, non-size increasing, program analysis, size change termination, termination

1. INTRODUCTION

1.1 Motivations

The goal of this study is an attempt to predict and control computational resources like space or time, which are used during the execution of a program. For this, we introduce a new tool called *Resource Control Graphs* and focus here on explaining how it can be used for termination proofs and space complexity management.

We present a data flow analysis of a low-level language sketched by means of Resource Control Graph, and we think that this is a generic concept from which several program properties could be checked.

Usual data flow analyses (see Nielson et al. [1999] for a detailed overview) use transfer functions to express how a given property is modified when following the program's execution. Then, a fixed point algorithm finds for each label a set of all possible values for the property. For example, one might be interested in which sign a given variable can take

Author's address: J.-Y. Moyen, LIPN, Institut Galilée, 99 avenue J.B. Clément, 93430 Villetaneuse, France.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1529-3785/YY/00-0001 \$5.00

at each point. The instructions of the program gives constraints on this (from one label to the next one). Iterating these constraints with a fixed point algorithm can find the set of all possible signs for the variable at each label.

Here, we want to consider each execution separately. So, when iterating the transfer function and coming back to an already treated label, instead of unifying of the new constraint with the old one and use fixed point, we will consider this as a new configuration. In the end, instead of having one set associated to each label, we will get a set of so called “walks”, each associating one value to each occurrence of each label. For example, a first walk can tell that if starting with a positive value at a given label, the variable will stay positive, but another walk tells that if starting with a negative value, the variable may become positive (while in such a case, the fixed point algorithm will build the set $\{+, -\}$ for each label).

Of course, we then need a way to study this set of walks and find common properties on them that tells something about the program.

The first problem we consider is the one of detecting programs able to compute within a constant amount of space, that is without performing dynamic memory allocation. These were dubbed *Non Size Increasing* by Hofmann [1999].

There are several approaches which try to solve this problem. The first protection mechanism is by monitoring computations. However, if the monitor is compiled with the program, it could itself cause memory leak or other problems. The second is the testing-based approach, which is complementary to static analysis. Indeed, testing provides a lower bound on the memory usage while static analysis gives an upper bound. The gap between both bounds is of some value in practice. Lastly, the third approach is type checking done by a bytecode verifier. In an untrusted environment (like embedded systems), the type protection policy (Java or .Net) does not allow dynamic allocation. Actually, the former approach relies on a high-level language, which captures and deals with memory allocation features [Aspinall and Compagnoni 2003]. Our approach guarantees, and even provides, a proof certificate of upper bound on space computation on a low-level language without disallowing dynamic memory allocations.

The second problem that we study is termination of programs. This is done by closely adapting ideas of Lee et al. [2001], Ben-Amram [2006] and Abel and Altenkirch [2002]. The intuition being that a program terminates whenever there is no more resources to consume.

There are long term theoretical motivations. Indeed a lot of work have been done in the last twenty years to provide syntactic characterisations of complexity classes, *e.g.* by Bellantoni and Cook [1992] or Leivant and Marion [1993]. Those characterisations are the bare bone of recent research on delineating broad classes of programs that run in some amount of time or space, like Hofmann, but also Niggel and Wunderlich [2006], Amadio et al. [2004], and Bonfante et al. [2007].

We believe that our Resource Control Graphs will be able to encompass several, or even all, of these analyses and express them in a common framework. In this sense, Resource Control Graphs are an attempt to build a generic tool for program analysis.

1.2 Coping with undecidability

All these theoretical frameworks share the common particularity of dealing with behaviours of programs (like time and space complexity) and not only with the inputs/outputs relation

which only depends on the computed function.

Indeed, a given function can be computed by several programs with different behaviours (in terms of complexity or other). Classical complexity theory deals with functions and considers *extensional* complexity. Here, we want to consider *intensional* complexity, that is try to understand why a given algorithm is more efficient than another to compute the same function.

The study of extensional complexity quickly reaches the boundary of Rice's theorem. Any extensional property of programs is either trivial or undecidable. Intuition and empirical results point out that intensional properties are even harder to decide.

However, several very successful works do exist for studying both extensional properties (like termination) or intensional ones (like time or space complexity). As these works provide decidable criteria, they must be either incomplete (reject a valid program) or unsound (accept an invalid program). Of course, the choice is usually to ensure soundness: if the program is accepted by the criterion, then the property (termination, polynomial bound, . . .) is guaranteed. This allows the criterion to be seen as a certificate in a proof carrying code paradigm.

When studying intensional properties, two different kinds of approaches exist. The first one consists of restricting the syntax of programs so that any program necessarily has the wanted property. This is in the line of the works on primitive recursive functions where the recurrence schemata is restricted to only primitive recursion. This approach gives many satisfactory results, such as the characterisations of PTIME by Cobham [1962] or Bellantoni and Cook [1992], the works of Leivant and Marion on tiering and predicative analysis [1993] or the works of Jones on CONS-free programs [2000]. On the logical side, this leads to explicit management of resources in Linear Logic [Girard 1987].

All these characterisations usually have the very nice property of *extensional completeness* in the sense that, *e.g.*, a function is in PTIME if and only if it can be defined by bounded primitive recursion (Cobham). Unfortunately, *intensionality* is not their main concern: these methods usually do not capture natural algorithms [Colson 1998], and programmers have to rewrite their programs in a non-natural way.

So, the motto of this first family of methods can be described as leaving the proof burden to the programmer rather than to the analyser. If one can write a program with the given syntax (which, in some cases, can be a real challenge), then certain properties are guaranteed. The other family of methods will go in the other way. Let the programmer write whatever he wants but the analysis is not guaranteed to work.

Since any program can *a priori* be given to the analysis, decidability is generally achieved by loosening the semantics during analysis. That is, one will consider *more* than all the executions a program can have. This approach is more recent but has already some very successful results such as the Size Change Termination [Lee et al. 2001] or the *mwp*-polynomials of Kristiansen and Jones [2005].

This second kind of methods can thus be described as not meddling with the programmer and let the whole proof burden lay on the analysis. Of course, the analysis being incomplete, one usually finds out that certain kinds of programs will not be analysed correctly and have to be rewritten. But this restriction is done *a posteriori* and not *a priori* and it can be tricky to find what exactly causes the analysis to fail.

Resource Control Graphs are intended to live within this second kind of analyses. Hence, the toy language used as an example is Turing-complete and will not be restricted.

1.3 Outline

Section 2 introduces the stack machines used all along as a simple yet powerful programming language. Section 3 describes the core idea of Resource Control Graphs that can be summed up as finding a decidable (recursive) superset of all the executions that still ensure a given property (such as termination or a complexity bound). Then, Section 4 immediately shows how this can be used in order to detect Non Size Increasing programs. Section 5 presents Vectors Addition Systems with States which are generalised into Resource Systems with States in Section 6. They form the backbone of the Resource Control Graphs. Section 7 present the tool itself and explain how to build a Resource Control Graph for a program and how it can be used to study the program. Section 8 shows application of RCG in building termination proofs similar to the Size Change Termination principle. Finally, Section 9 discuss how matrices multiplication is or could be used in program analyses thus leading to several possible further developments.

1.4 Notations

In a directed graph¹ $G = (S, A)$, will write $s \xrightarrow{a} s'$ to say that a is an edge between s and s' . Similarly, we will write $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n$ to say that $a_1 \dots a_n$ is a path going through vertices s_0, \dots, s_n . Or simply $s_0 \xrightarrow{w} s_n$ if $w = a_1 \dots a_n$. $s \rightarrow s'$ means that there exists an edge a such that $s \xrightarrow{a} s'$ and $\xrightarrow{+}, \xrightarrow{*}$ are the transitive and reflexive-transitive closures of \rightarrow .

A partial order \prec is a well partial order if there are no infinite anti-chain, that is for every infinite sequence x_1, \dots, x_n, \dots there are indexes $i < j$ such that $x_i \preceq x_j$. This mean that the order is well-founded (no infinite decreasing sequence) but also that there is no infinite sequence of pairwise incomparable elements. The order induced by the divisibility relation on \mathbb{N} , for example, is well-founded but is not a well partial order since the sequence of all prime numbers is an infinite sequence of pairwise incomparable elements.

The set of integers (positive and negative) is \mathbb{Z} , and \mathbb{N} is the set of integers ≥ 0 , when working with infinity, $\overline{\mathbb{Z}} = \mathbb{Z} \cup \{+\infty\}$, that is we do not need $-\infty$ here. When working with vectors of \mathbb{Z}^k , \leq denotes the component-wise partial order. That is $a \leq b$ if and only if $a_i \leq b_i$ for all $1 \leq i \leq k$. This is a well partial order on \mathbb{N}^k .

2. STACK MACHINES

2.1 Syntax

A stack machine consist of a finite number of *registers*, each able to store a letter of an alphabet, and a finite number of *stacks*, that can be seen as lists of letters. Stacks can only be modified by usual `push` and `pop` operations, while registers can be modified by a given set of operators each of them assumed to be computed in a single unit of time.

¹We will use $s \in S$ to designate vertices and $a \in A$ to designates edges. The choice of using French initials (“Sommet” and “Arête”) rather than the usual (V, E) is done to avoid confusion between vertices and the valuations introduced later.

Definition 2.1 (Stack machine). Stack machines are defined by the following grammar:

(Alphabet)	Σ	finite set of symbols
(Programs)		$p ::= \text{lbl}_1 : i_1; \dots; \text{lbl}_n : i_n;$
(Instructions)	$\mathcal{I} \ni i ::=$	$\text{if (test) then goto lbl}_0 \text{ else goto lbl}_1 \mid$ $\mathbf{r} := \text{pop}(\text{stk}) \mid \text{push}(\mathbf{r}, \text{stk}) \mid \mathbf{r} := \text{op}(\mathbf{r}_1, \dots, \mathbf{r}_k) \mid \text{end}$
(Labels)	$\mathcal{L} \ni \text{lbl}$	finite set of labels
(Registers)	$\mathcal{R} \ni \mathbf{r}$	finite set of registers
(Stacks)	$\mathcal{S} \ni \text{stk}$	finite set of stacks
(Operators)	$\mathcal{O} \ni \text{op}$	finite set of operators

Each operator has a fixed arity k and n is an integer constant. The syntax of a program induces a function $\text{next} : \mathcal{L} \rightarrow \mathcal{L}$ such that $\text{next}(\text{lbl}_i) = \text{lbl}_{i+1}$ and a mapping $\iota : \mathcal{L} \rightarrow \mathcal{I}$ such that $\iota(\text{lbl}_k) = i_k$. The `pop` operation removes the top symbol of a stack and put it in a register. The `push` operation copies the symbol in the register onto the top of the stack. The `if`-instruction gives control to either `lbl0` or `lbl1` depending on the outcome of the test. Each operator is interpreted with respect to a given semantics function $\llbracket \text{op} \rrbracket$.

The precise sets of labels, registers and stacks can be inferred from the program. Hence if the alphabet is fixed, the machine can be identified with the program itself.

The syntax `lbl : if (test) then goto lbl0` can be used as a shorthand for `lbl : if (test) then goto lbl0 else goto next(lbl)`. Similarly, we can abbreviate `if true then goto lbl` as `goto lbl`, that is an unconditional jump to a given label. What kind of tests can be used is not specified here. Of course, tests must be computable (for obvious reasons) in constant time and space (so that they do not play an important part when dealing with complexity properties). Comparisons between letters of the alphabet (e.g. \leq if they are integers) are typical tests that can be used.

If the alphabet contains a single letter, then the registers are useless and the stacks can be seen as unary numbers. The machine then becomes an usual counter machine [Shepherdson and Sturgis 1963].

Example 2.2. The following program reverses a list in stack `l` and put the result in stack `l'`. It uses register `a` to store intermediate letters. The empty stack is denoted `[]`.

0 : if <code>l = []</code> then goto end;	3 : goto 0;
1 : <code>a := pop(l)</code> ;	end : end;
2 : <code>push(a, l')</code> ;	

2.2 Semantics

Definition 2.3 (Stores). A *store* is a function σ assigning a symbol (letter of the alphabet) to each register and a finite string in Σ^* to each stack. Store update is denoted $\sigma\{x \leftarrow v\}$.

Definition 2.4 (States). Let p be a stack program. A *state* of p is a couple $\theta = \langle \text{IP}, \sigma \rangle$ where the *Instruction Pointer* `IP` is a label and σ is a store. Let Θ be set of all states, Θ^* (Θ^ω) be the set of finite (infinite) sequences of states and $\Theta^{*\omega}$ be the union of both.

$$\begin{array}{c}
\frac{i = \iota(\text{IP}) = \mathbf{r} := \text{op}(\mathbf{r}_1, \dots, \mathbf{r}_k) \quad \sigma' = \sigma\{\mathbf{r} \leftarrow \llbracket \text{op} \rrbracket(\sigma(\mathbf{r}_1), \dots, \sigma(\mathbf{r}_k))\}}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{i} \langle \text{next}(\text{IP}), \sigma' \rangle} \\
\\
\frac{\iota(\text{IP}) = \text{if } (\text{test}) \text{ then goto } \text{lbl}_1 \text{ else goto } \text{lbl}_2 \quad (\text{test}) \text{ is true}}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{(\text{test})_{\text{true}}} \langle \text{lbl}_1, \sigma \rangle} \\
\\
\frac{\iota(\text{IP}) = \text{if } (\text{test}) \text{ then goto } \text{lbl}_1 \text{ else goto } \text{lbl}_2 \quad (\text{test}) \text{ is false}}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{(\text{test})_{\text{false}}} \langle \text{lbl}_2, \sigma \rangle} \\
\\
\frac{i = \iota(\text{IP}) = \mathbf{r} := \text{pop}(\text{stk}) \quad \sigma(\text{stk}) = \lambda.w \quad \sigma' = \sigma\{\mathbf{r} \leftarrow \lambda, \text{stk} \leftarrow w\}}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{i} \langle \text{next}(\text{IP}), \sigma' \rangle} \\
\\
\frac{i = \iota(\text{IP}) = \mathbf{r} := \text{pop}(\text{stk}) \quad \sigma(\text{stk}) = \epsilon}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{i} \perp} \\
\\
\frac{i = \iota(\text{IP}) = \text{push}(\mathbf{r}, \text{stk}) \quad \sigma' = \sigma\{\text{stk} \leftarrow \sigma(\mathbf{r}).\sigma(\text{stk})\}}{p \vdash \langle \text{IP}, \sigma \rangle \xrightarrow{i} \langle \text{next}(\text{IP}), \sigma' \rangle}
\end{array}$$

Fig. 1. Small steps semantics

Definition 2.5 (Executions). The operational semantics of Figure 1 defines a relation² $p \vdash \theta \xrightarrow{i} \theta'$.

An *execution* of a program p is a sequence (finite or not) $p \vdash \theta_0 \xrightarrow{i_1} \theta_1 \xrightarrow{i_2} \dots \xrightarrow{i_n} \theta_n \dots$.

An infinite execution is said to be *non-terminating*. A finite execution is *terminating*. If the program admits no infinite execution, then it is *uniformly terminating*.

We use \perp to denote runtime error. We may also allow operators to return \perp if we want to allow operators generating errors. It is important to notice that \perp is not a state, and hence, will not be considered when quantifying over all states.

If the instruction is not specified, we will write simply $p \vdash \theta \rightarrow \theta'$ and use $\xrightarrow{+}$, $\xrightarrow{*}$ for the transitive and reflexive-transitive closures.

Definition 2.6 (Traces). The *trace* of an execution $p \vdash \theta_0 \xrightarrow{i_1} \theta_1 \xrightarrow{i_2} \dots \xrightarrow{i_n} \theta_n \dots$ is the instructions sequence $i_1 \dots i_n \dots$.

Definition 2.7 (Length). Let $\theta = \langle \text{IP}, \sigma \rangle$ be a state. Its *length* $|\theta|$ is the sum of the number of elements in each stack³. That is:

$$|\theta| = \sum_{\text{stk} \in \mathcal{S}} |\text{stk}|$$

²Notice that the label i on the edge is technically not an instruction since for tests we also keep the information of which branch is taken.

³Hence, it should more formally be $|\langle \text{IP}, \sigma \rangle| = \sum_{\text{stk}_i \in \mathcal{S}} |\sigma(\text{stk}_i)|$. Since explicitly mentioning the store everywhere would be quite unreadable, we use stk_i instead of $\sigma(\text{stk}_i)$ and, similarly, \mathbf{r} instead of $\sigma(\mathbf{r})$, when the context is clear.

Length is the usual notion of space. Since there is a fixed number of registers and each can only store a finite number of different values, the space need to store all registers is always bounded. So, we do not take registers into account while computing space usage.

The notion of length allows to define usual time and space complexity classes.

Definition 2.8 (Running time, running space). The *time usage* of a finite execution is the number of states in it. Let f be an increasing function from the non-negative integers to the non-negative integers. We say that the *running-time* of a program is bounded by f if the time usage of each execution is bounded by $f(|\theta|)$ where θ is the first state of the execution.

The *space usage* of a finite execution is the maximum length of a state in it. Let f be an increasing function from the non-negative integers to the non-negative integers. We say that the *running-space* of a program is bounded by f if the space usage of each execution is bounded by $f(|\theta|)$ where θ is the first state of the execution.

Definition 2.9 (Complexity). Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be an increasing function. The class $T(f)$ is the set of functions which can be computed by a program whose running time is bounded by f . The class $S(f)$ is the set of function which can be computed by a program whose running space is bounded by f .

As usual, PTIME denotes the set of all functions computable in polynomial time, that is the union of $T(P)$ for all polynomials P and so on.

If we want to define classes such as LOGSPACE, then we must, as usual, use some read-only stacks which can only be popped but not pushed and some write-only output stacks. These play no role when computing the length of a state.

2.3 Turing Machines

Stack machines are Turing complete. We quickly describe here the straightforward way to simulate a Turing machine by a stack machine.

Simulating a TM with a single tape and alphabet Σ can be done with a stack machine with the alphabet $\Sigma \cup Q$ (where Q is the set of states of the TM), two stacks and two registers. The two stacks and the first register will encode the tape in an usual way (one stack, reversed, for the left-hand side, the register for the scanned symbol and the other stack for the right-hand side). Another register will contain the current state of the automaton.

At each step, the program will go through a sequence of tests on the state and scanned symbol in order to find the right set of instructions to perform and after that jump back to the beginning of the program. There will be at most $q \times m$ such tests where q is the number of states of the TM and m the number of symbols in the alphabet. Then, simulation of a step is quite easily done by modifying the “scanned symbol” register and then simulating movement.

Simulating movement first has to check that the correct stack is not empty, push a “blank symbol” on it if necessary and then push the scanned symbol on one stack and pop the other stack onto it.

Each step of the TM is simulated in a constant number of steps of the stack machine (depending only on the TM). So that the time complexity of the stack machine will be the same as the time complexity of the TM (up to a multiplicative constant). Similarly, at any step of the simulation, the length of the configuration of the stack machine will be the number of non-blank or scanned symbols on the tape (minus one because one symbol is stored into a register). So the space complexity will be the same.




Fig. 2. Sequences of states, executions and admissible sequences

3. A TASTE OF RCG

This section describes the idea behind Resource Control Graphs in order to get a better grip on the formal definitions later on.

3.1 Admissible sequences

Consider an execution of a program. It can be described as a sequence of states. Clearly, not all sequences of states describe an execution. So we have a set of executions, Υ , which is a subset of the set of all sequences of states (finite or infinite), $\Theta^{*\omega}$.

The undecidability results mean that given a program it is impossible to say if the set of executions, Υ , and Θ^ω , the set of infinite sequences of states, are disjoint. So, the idea here is to find a set \mathcal{A} of *admissible* sequences, which is a superset of the set of all executions, and whose intersection with Θ^ω can be computed. If this intersection is empty, then *a fortiori*, there are no infinite executions of the program, but if the intersection is not empty, then we cannot decide if this is due to some non-terminating execution of the program or to some of the sequences added for the sake of the analysis. This means that depending on the machine considered and the way \mathcal{A} is build, we can be in three different situations as depicted in Figure 2. We build $\mathcal{A} \supset \Upsilon$ such that $\mathcal{A} \cap \Theta^\omega$ is decidable. If it is empty, then the program uniformly terminates; otherwise, we cannot say anything. Of course, the undecidability theorem means that if we require \mathcal{A} to be recursive (or at least recursively separable from Θ^ω), then there will necessarily be some programs for which the situation will be the one in the middle (in Figure 2), that is we falsely suppose that the program does not uniformly terminate.

One convenient way to represent all the possible executions (and only these), is to build a *state-transition graph*. This is a directed graph where each vertex is a state of the program and there is an edge between two vertices if and only if it is possible to go from one state to the other with a single step of the operational semantics. Of course, since there are infinitely many different stores, there are infinitely many possible states and the graph is infinite.

3.2 The folding trick

Using the state-transition graph to represent executions is not convenient since handling an infinite graph can be tedious. To circumvent this, we must look into states and decompose them.

A state is actually a couple of one label and one store. The label corresponds to the *control* of the program while the store represents *memory*. A first try to get ride of the infinite state-transition graph is then to only consider the control part of each state.

Thus, there will only be finitely many different nodes in the graph (since there are only finitely many different labels). By identifying all states bearing the same label, it becomes




Fig. 3. CFG of the reverse program.

possible to “fold” the infinite state-transition graph into a finite graph, called the *Control Flow Graph* (CFG) of the program. The CFG is an usual tool for program analyses and transformations and can directly be built from the program.

Definition 3.1 (Control Flow Graph). Let p be a program. Its *Control Flow Graph* (CFG) is a directed graph $G = (S, A)$ where:

- $S = \mathcal{L}$. There is one vertex for each label.
- If $\iota(\text{lbl}) = \text{if } (\text{test}) \text{ then goto } \text{lbl}_1 \text{ else goto } \text{lbl}_2$ then there is one edge from lbl to lbl_1 labelled $(\text{test})_{\text{true}}$ and one from lbl to lbl_2 labelled $(\text{test})_{\text{false}}$.
- If $\iota(\text{lbl}) = \text{end}$ then there is no edge going out of lbl .
- Otherwise, there is one edge from lbl to $\text{next}(\text{lbl})$ labelled $\iota(\text{lbl})$.

Vertices and edges are named after, respectively, the label or instruction⁴ they represent. No distinction is made between the vertex and the label or the edge and the instruction as long as the context is clear.

Example 3.2. The CFG of the reverse program is displayed on Figure 3.

With state-transition graphs, there was a one-to-one correspondence between executions of the program and (maximal) paths in the graph. This is no longer true with Control Flow Graphs. Now, to each execution corresponds a path (finite or infinite) in the CFG. The converse, however, is not true. There are paths in the CFG that correspond to no execution.

Let \mathcal{P} be the set of paths in the CFG. \mathcal{P} is a regular language over the alphabet of the edges (see Lemma 5.9), hence \mathcal{P} is recursive. Since we can associate a path to each execution, we can say that \mathcal{P} is a superset of Υ .

This leads to a first try at building a set of admissible sequences by choosing $\mathcal{A} = \mathcal{P}$.

However, as soon as the graph contains loops, \mathcal{P} will contain infinite sequences. So this is quite a poor try at building an admissible set of sequences, corresponding exactly to the trivial analysis “A program without loop uniformly terminates”.

In order to do better, we need to plug back the memory into the CFG.

⁴Again, since the two branches of tests are separated, some edges do not correspond exactly to an instruction of the program. We will nonetheless continue to call these “instructions”.

3.3 Walks

So, in order to take memory into account but still keep the CFG, we will not consider vertices any more but states again. Clearly, each state is associated to a vertex of the CFG. Moreover to each instruction i , we can associate a function $\llbracket i \rrbracket$ such that for all states θ, θ' such that $p \vdash \theta = \langle \text{IP}, \sigma \rangle \xrightarrow{i} \langle \text{IP}', \sigma' \rangle = \theta'$, we have $\sigma' = \llbracket i \rrbracket(\sigma)$.

So, instead of considering paths in the graph, we can now consider walks. Walks are sequences of states following a path where each new store is computed according to the semantics function $\llbracket i \rrbracket$ of the edge just followed.

The only case where the CFG has out-degree greater than 1 is for tests. In order to prevent the wrong branch to be taken, the semantics function $\llbracket (\text{test})_{\text{true}} \rrbracket$ can be a partial function only defined for stores where the test is true (and conversely for the false branch of tests).

But if we do this exactly that way, then there will be a bijection between the executions and the walks and everything will stay undecidable.

So the idea at this point is to keep both branches of the test possible, that is more or less replacing a deterministic test by a non-deterministic choice between the two outcomes. This leads to a set of walks bigger than the set of executions but, hopefully, recursively separable from the set of infinite sequences of states.

4. MONITORING SPACE USAGE

In order to illustrate the ideas of previous Section, we introduce here the notion of Resource Control Graph for the specific case of monitoring space usage. In Section 7, this notion will be fully generalised to define Resource Control Graphs.

4.1 Space Resource Control Graphs

Definition 4.1 (Weight). For each instruction i , we define a *weight* k_i as follows:

- The weight of any instruction that is neither push nor pop is 0.
- The weight of a push instruction is +1.
- The weight of a pop instruction is -1.

PROPOSITION 4.2. For all states θ such that $p \vdash \theta \xrightarrow{i} \theta'$, we have $|\theta'| = |\theta| + k_i$.

It is important here that both θ and θ' are states. Indeed, this means that when an error occurs (\perp), we remove all constraints.

Definition 4.3 (Space Resource Control Graph). Let p be a program. Its Space Resource Control Graph (Space-RCG) is a weighted directed graph G such that:

- G is the Control Flow Graph of p .
- For each edge i , the weight $\omega(i)$ is k_i .

Definition 4.4 (Configurations, walks). A *configuration* is a couple $\eta = (s, v)$ where $s \in S$ is a vertex and $v \in \mathbb{Z}$ is the *valuation*. A configuration is *admissible* if and only if $v \in \mathbb{N}$.

A *walk* is a sequence (finite or not) of configurations $(s_0, v_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, v_n) \xrightarrow{a_{n+1}} \dots$ such that $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n \xrightarrow{a_{n+1}} \dots$ and for all $i > 0$, $v_i = v_{i-1} + \omega(a_i)$. A walk is *admissible* if all configurations in it are admissible.

Definition 4.5 (Traces). The *trace* of a walk is the sequence of all edges followed by the walk, in order.

PROPOSITION 4.6. *Let p be a program, G be its Space-RCG and $p \vdash \theta_1 = \langle \mathcal{IP}_1, \sigma_1 \rangle \rightarrow \dots \rightarrow \theta_n = \langle \mathcal{IP}_n, \sigma_n \rangle$ be an execution with trace t , then there is an admissible walk in G , $(\mathcal{IP}_1, |\theta_1|) \rightarrow \dots \rightarrow (\mathcal{IP}_n, |\theta_n|)$, with the same trace t .*

PROOF. By construction of the Space-RCG and induction on the length of the execution. \square

4.2 Characterisation of Space usage

THEOREM 4.7. *Let f be a total function $\mathbb{N} \rightarrow \mathbb{N}$. Let p be a program and G be its Space-RCG.*

$p \in S(f)$ if and only if for each state $\theta_0 = \langle \mathcal{IP}_0, \sigma_0 \rangle$ and each execution $p \vdash \theta_0 \xrightarrow{} \theta_n$, the trace of the execution is also the trace of an admissible walk $(\mathcal{IP}_0, |\theta_0|) \rightarrow (\mathcal{IP}_1, i_1) \rightarrow \dots \rightarrow (\mathcal{IP}_n, i_n)$ and for each k , $i_k \leq f(|\theta_0|)$.*

PROOF. Proposition 4.6 tells us that $i_k = |\theta_k|$. Then, both implications hold by definition of space usage. \square

Definition 4.8 (Resource awareness). A Space-RCG is *f -resource aware* if for any walk $(s_0, v_0) \xrightarrow{*} (s_n, v_n)$, $v_n \leq f(v_0)$.

COROLLARY 4.9. *Let $f : \mathbb{N} \rightarrow \mathbb{N}$ be a total function, p be a program and G be its Space-RCG.*

If G is f -resource aware, then $p \in S(f)$.

Here, the converse is not true because the Space-RCG can have admissible walks with uncontrolled valuations which do not correspond to any real execution.

4.3 Non Size Increasingness

The study of Non Size Increasing (NSI) functions was introduced by Hofmann [1999]. Former syntactical restrictions for PTIME, such as the safe recurrence of Bellantoni and Cook [1992], forbid to iterate a function because this can yield to super-polynomial growth. However, this prevents from using perfectly regular algorithms such as the insertion sort where the insertion function is iterated. The idea is then that iterating functions *who do not increase the size of data* is harmless.

In order to detect these functions, Hofmann uses a typed functional programming language. A special type, \diamond , is added. There are no constructors for this type, meaning that the only normal forms of type \diamond are variables. It can be seen as the type of pointers to free memory. Constructors of other types now require one or more \diamond . Typically, the usual `cons` for lists requires a \diamond in addition to the data and the list and will then be typed $\text{cons} : \diamond \times \alpha \times L(\alpha) \rightarrow L(\alpha)$. When building a list, the \diamond must be a variable (no closed term of type \diamond), similar to a pointer to a free cell in memory where the list can be built.

Whenever a list is destroyed, the \diamond in the `cons` is freed (in a variable) and can thus be later reused to build another list. By ensuring a linear type discipline, one can be sure that no \diamond is ever duplicated. Then, any program that can be typed with this type system can be computed in a NSI way, e.g. be compiled into C without any `malloc` instruction.

Example 4.10. Here is the insertion sort “*the NSI-way*”. Notice how the \diamond s freed by the pattern matching are then reused when the resulting list is built. With usual chained

lists in \mathbb{C} (struct with a value and a pointer to next cell), this corresponds to the idea of reusing the same pointers over and over rather than actually freeing memory at pattern matching before reallocating it afterwards.

```

insert :  $\diamond \times \alpha \times \text{List}(\alpha) \rightarrow \text{List}(\alpha)$ 
      insert( $d, a, \mathbf{nil}$ )  $\rightarrow$  cons( $d, a, \mathbf{nil}$ )
insert( $d, a, \text{cons}(d', b, l)$ )  $\rightarrow$  if  $a < b$ 
      then cons( $d, a, \text{cons}(d', b, l)$ )
      else cons( $d', b, \text{insert}(d, a, \mathbf{nil})$ )

sort :  $\text{List}(\alpha) \rightarrow \text{List}(\alpha)$ 
      sort( $\mathbf{nil}$ )  $\rightarrow$   $\mathbf{nil}$ 
      sort( $\text{cons}(d, a, l)$ )  $\rightarrow$  insert( $d, a, \text{sort}(l)$ )

```

With Space RCG, valuations in a walk play the same role as Hofmann's diamonds (\diamond). The higher the value, the more diamonds are needed in the current configuration. `push` has positive weight, meaning that it uses diamonds but `pop` has negative weight, meaning that it releases diamonds for later use.

Definition 4.11 (Non Size Increasing). A program is *Non Size Increasing* (NSI) if its space usage is bounded by $\lambda x.x + \alpha$ for some constant α .

NSI is the class of functions which can be computed by Non Size Increasing programs. That is $\bigcup_{\alpha} S(\lambda x.x + \alpha)$.

PROPOSITION 4.12. *Let p be a program and G be its Space-RCG. If G is $\lambda x.x + \alpha$ -resource aware for some constant α , then p is NSI.*

PROOF. This is a direct consequence of Theorem 4.7. \square

THEOREM 4.13. *Let p be a program and G be its Space-RCG. G is $\lambda x.x + \alpha$ -resource aware (for some α) if and only if it contains no cycle of strictly positive weight.*

PROOF. If there is no cycle of strictly positive weight, then let α be the maximum weight of any path in G . Since there is no cycle of strictly positive weight, it is well-defined. Consider a walk $(s_0, v_0) \xrightarrow{*} (s_n, v_n)$ in G . Since α is the maximum weight of a path, we have $v_n \leq v_0 + \alpha$. Hence, G is $\lambda x.x + \alpha$ -resource aware.

Conversely, if there is a cycle of strictly positive weight, then it can be followed infinitely many times and provides an admissible walk with unbounded valuations. \square

Building the Space-RCG can be done in linear time in the size of the program. Finding the maximum weight of a path can be done in polynomial time in the size of the graph (so in the size of the program) with Bellman-Ford's algorithm ([Cormen et al. 1990] chapter 25.5). So we can detect NSI programs and find the constant α in polynomial time in the size of the program.

Example 4.14. The Space-RCG of the reverse program (from Example 2.2) is displayed on Figure 4. Since it contains no cycle of strictly positive weight, the program




Fig. 4. Space-RCG of the reverse program.

is Non Size Increasing. Moreover, since the maximum weight of any path is 0, it can be computed in space $\lambda x.x$, that is the constant α is 0 for this program.

This result, however, lacks an intensionality statement (how much of all NSI programs are caught?) or even an extensional completeness one (does there exist functions in NSI that are not captured by such a program?) Of course, the class of all *programs* which are Non Size Increasing is undecidable. This means that intensionality statements are hard to achieve. However, we can reach an extensional completeness one.

Without loss of generality, we consider here that in the initial configuration of a TM, the tape consists in only blank symbols except for a *consecutive* sequence of symbols which are *all* non-blank. That is, we do not allow input tape to have the shape $x \sqcup y$ where x and y are non-blank symbols and \sqcup is the blank symbol. This allows to detect the end of input as the first blank symbol. The head is assumed to scan the leftmost non-blank symbol at the beginning of computation.

PROPOSITION 4.15 (NORMALISING TMS). *Let M be a NSI Turing Machine running in space $\lambda x.x + \alpha$. There exists a TM \widetilde{M} , computing the same function, running in space $\lambda x.x + \alpha + 2$, proceeding in 2 phases:*

- (1) *Firstly, \widetilde{M} writes 2 $\#$ and α B on blank squares of its tape, where both $\#$ and B are new symbols.*
- (2) *Secondly, \widetilde{M} never scans a blank symbol again.*

PROOF. \widetilde{M} starts by going one square left and writing $\#$ there. Then it goes to the end of the input, write α B after it (since α is fixed for M and does not depend on the input, this can be done) and finally another $\#$. Then, it goes back to the beginning of the input (the symbol immediately after the $\#$) and goes into the second phase.

In the second phase, \widetilde{M} simulates M . However, $\#$ are never overwritten. Whenever M request to write a $\#$, the whole content of the tape is shifted one square left (or right) and simulation of M resumes where it stopped. Since M is NSI, the simulation can be done entirely between the two $\#$. \square

Of course, this simulation is rather costly from a time point of view, but since we are only concerned with space here, that does not matter. Notice also that such a normalisation could be made for a TM running in space $f(x)$ for any computable function f . However, in that case the simulation would require an additional tape to compute $f(x)$ from the input

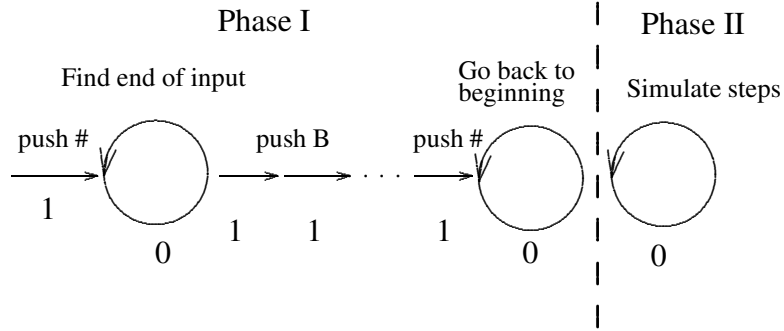


Fig. 5. Space-RCG of the simulation of a normalised NSI Turing machine.

and then allocate sufficiently many new squares. This would be quite tricky to do and require control over the space used to compute $f(x)$...

THEOREM 4.16 (EXTENSIONAL COMPLETENESS). *Let M be a NSI TM, \tilde{M} be the corresponding normalised machine and \tilde{p} be the program simulating \tilde{M} according to the simulation of Section 2.3. Let \tilde{G} be the Space-RCG of \tilde{p} .*

\tilde{G} contains no cycle of strictly positive weight.

PROOF. During the second phase of the simulation of M , \tilde{M} never scan a blank symbol. Hence, there is no need to push new (blank) symbols on any of the stacks. While moving the head, each push on one stack is immediately followed by a pop on the other stack, thus yielding only paths of weight 0.

During the first phase of the simulation, p starts by adding a symbol on a stack (a blank symbol immediately erased by #, or alternatively directly a # with a slightly smarter simulation). Then it loops to find the end of the input. During this loop, each push is also followed by a pop, thus creating only cycles of weight 0. Then it adds $\alpha + 1$ new symbols (B and #), but since α does not depend on the input, this can be done by $\alpha + 1$ separate push, thus creating no cycles. And lastly it goes back to the start of the input, again each push is followed immediately by a pop. Figure 5 shows how the Space-RCG of \tilde{p} looks like. \square

This result means that our characterisation of NSI is extensionally complete. Each function in NSI can be computed by a program which fits into the characterisation (that is, whose Space-RCG is $\lambda x. x + \alpha$ -resource aware). Of course, intentional completeness (capturing all NSI programs) is far from reached (but is unreachable with a decidable algorithm).

4.4 Linear Space

LINSPACE seems to be closely related to NSI. Indeed, LINSPACE functions can be computed in space $\lambda x. \beta x + \alpha$ and so NSI is a special case of LINSPACE with $\beta = 1$. So we want to adapt our result to detect linear space usage.

The idea is quite easy: since we are allowed to use β time more space than what is initially allocated, it is sufficient to consider that every time some of the initial data is freed, β “tokens” (\diamond) are released and can later be used to control β different allocations.

In order to do so, the most convenient way is to design certain stacks of the machine (or

certain tapes of a TM) as *input stacks* and the others must be initially empty. Then, a `pop` operation over an input stack would have weight $-\beta$ instead of simply -1 to account for this linear factor. However, doing so we must be careful that newly allocated memory (that is, further `push`) will only be counted as 1 when freed again (to avoid a cycle of freeing one slot, allocating β , freeing these β slots and reallocating β^2 and so on). In order to do so, we simply require that the input stacks are read-only in the sense that it is not possible to perform a `push` operation on them.

Notice that any program can be turned into such kind of program by having twice as many stacks (one input and one work for each) and starting by copying all the input stacks into the corresponding working stacks and then only dealing with the working stacks.

With these programs, the invariant will not be the length of states, but something slightly more complicated, namely β times the length of input stacks plus the length of work stacks. We will call this measure *size*. Globally, we will use *size* to denote some kind of measure on states that is used by the RCG for analysis. The terminology is close from the one of the Size Change Termination [Lee et al. 2001] where values are assumed to have some (well-founded) “size ordering” which is not specified and not necessarily related to the actual space usage of the data. Typically, termination of a program working over positive integers can be proved using the usual ordering on \mathbb{N} as size ordering, even if the integers are all 32 bits integers, thus taking exactly the same space in memory.

Definition 4.17 (Extended stack machines). An *extended stack machine* is a stack machine with the following modification:

There are two disjoint sets of stacks, \mathcal{S}_i is the set of *input stacks* and \mathcal{S}_w is the set of *working stacks*. There are two instructions `popi` and `popw` depending on whether an input or working stack is considered but only one `push = pushw` instruction, that is it is impossible to `push` anything on an input stack.

The β -*size* of a state is β times the length of input stacks plus the length of working stacks, that is:

$$\|\theta\|_\beta = \beta \sum_{stk_i \in \mathcal{S}_i} |stk_i| + \sum_{stk_w \in \mathcal{S}_w} |stk_w|$$

The *weight* of `popi` is $-\beta$, the weight of `popw` is -1 , the weight of `push` is $+1$. the weight of other instructions is 0.

The β -Space RCG is build as the Space-RCG: the underlying graph is the control flow graph and the weight of each edge is the weight of the corresponding instruction.

Proposition 4.6 becomes:


PROPOSITION 4.18. *Let p be a program, G_β be its β -Space RCG and $p \vdash \theta_1 = \langle IP_1, \sigma_1 \rangle \rightarrow \dots \rightarrow \theta_n = \langle IP_n, \sigma_n \rangle$ be an execution with trace t , then there is an admissible walk $(IP_1, \|\theta_1\|_\beta) \rightarrow \dots \rightarrow (IP_n, \|\theta_n\|_\beta)$ with the same trace t .*

Then, adapting Theorem 4.7 and Theorem 4.13, we have:

PROPOSITION 4.19. *Let p be a program and G_β be its β -Space RCG. If G_β is $\lambda x.x + \alpha$ -resource aware for some constant α , then $S(p) \leq \lambda x.\beta x + \alpha$.*

THEOREM 4.20. *Let p be a program and G_β be its β -Space RCG. G_β is $\lambda x.x + \alpha$ -resource aware (for some α) if and only if it contains no cycle of strictly positive weight.*

COROLLARY 4.21. *Let p be a program. If there exists β such that its β -Space RCG contains no cycle of strictly positive weight, then p is in LINSPEACE.*

Fig. 6. β -Space RCG of the double-reverse program.

This can be checked in NPTIME since β is polynomially bounded in the size of the program.

For LINSPEACE also, the normalisation process of Turing Machine can quite easily be performed, typically by using an input (read-only) tape and a working tape where space usage is counted. The first phase of the normalised TM consist in repeatedly copy one symbol from the input tape to the right of the working tape and add $\beta - 1$ B at the left of the working tape, then putting the two $\#$ on the working tape. This means that here also the characterisation is extensionally complete: for each LINSPEACE function, there exists one program computing it that fits into the characterisation.

Example 4.22. The following program “double-reverses” a list. It is similar to the reverse program but each element is present twice in the result. The list \mathbf{l} is an input stack (and hence cannot be pushed) while \mathbf{l}' is a working stack.

0 : if $\mathbf{l} = []$ then goto end;	3 : push _w (\mathbf{a}, \mathbf{l}');
1 : $\mathbf{a} := \text{pop}_i(\mathbf{l})$;	4 : goto 0;
2 : push _w (\mathbf{a}, \mathbf{l}');	end : end;

Its β -Space RCG is displayed on Figure 6. Since it contains no cycle of strictly positive weight if $\beta \geq 2$, the program is in LINSPEACE, more precisely, it can be computed in space $\lambda x.2x$

5. VECTOR ADDITION SYSTEM WITH STATES

This section describes Vectors Addition Systems with States (VASS). Resources Control Graphs are a generalisation of VASS. VASS are known to be equivalent to Petri Nets [Reutenauer 1989].

5.1 Definitions

Definition 5.1 (VASS, configurations, walks). A *Vector Addition System with States* is a directed graph $G = (S, A)$ together with a *weighting function* $\omega : A \rightarrow \mathbb{Z}^k$ where k is a fixed integer.

A *configuration* is a couple $\eta = (s, v)$ where $s \in S$ is a vertex and $v \in \mathbb{Z}^k$ is the *valuation*. A configuration is *admissible* if and only if $v \in \mathbb{N}^k$.

A *walk* is a sequence (finite or not) of configurations $(s_0, v_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, v_n)$ such that $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n$ and for all $i > 0$, $v_i = v_{i-1} + \omega(a_i)$. A walk is *admissible* if all configurations in it are admissible.

We say that path $a_1 \dots a_n$ is the *underlying path* of the walk and the walk *follows* this path. Similarly, G is the *underlying graph* for the VASS.

As for graphs and paths, we will write $\eta \rightarrow \eta'$ if there exists an edge a such that $\eta \xrightarrow{a} \eta'$ and $\xrightarrow{+}, \xrightarrow{*}$ for the closures.

Definition 5.2 (Weight of a path). Let V be a VASS and $a_1 \dots a_n$ be a path in it. The weight of edges is extended to paths canonically: $\omega(a_1 \dots a_n) = \sum \omega(a_i)$. This means that ω is a morphism between (A, \cdot) (the free monoid generated by the edges) and $(\mathbb{Z}^k, +)$.

LEMMA 5.3. *Let V be a VASS and $a_1 \dots a_n$ be a finite path in it. There exists a valuation v_0 such that for $0 \leq i \leq n$, $v_0 + \omega(a_1 \dots a_i) \in \mathbb{N}^k$.*

This means that every finite path is the underlying path of an admissible walk.

PROOF. Because the path is finite, the j th component of $\omega(a_1 \dots a_i)$ is bounded from below by some α_j (of course, this bound is not necessarily reached with the same i for all components, but nonetheless such a bound exists for each component separately). By putting $\beta_j = \max(0, -\alpha_j)$ (that is 0 if α_j is positive), then $v_0 = (\beta_1, \dots, \beta_k)$ verifies the property. \square

LEMMA 5.4. *Let $(s_0, v_0) \rightarrow \dots \rightarrow (s_n, v_n)$ be an admissible walk in a VASS. Then, for all $v'_0 \geq v_0$ (component-wise comparison), $(s_0, v'_0) \rightarrow \dots \rightarrow (s_n, v'_n)$ is an admissible walk (following the same path).*

PROOF. By monotonicity of the addition. \square

Definition 5.5 (Uniform termination). A VASS is said to be *uniformly terminating* if it admits no infinite admissible walk. That is, every walk is either finite or reaches a non-admissible configuration.

THEOREM 5.6. *A VASS is not uniformly terminating if and only if there exists a cycle whose weight is in \mathbb{N}^k (that is, is non-negative with respect to each component).*

PROOF. If such a cycle exists, starting and ending at vertex s , then by Lemma 5.3 there exists v_0 such that the walk starting at (s, v_0) and following it is admissible. After following the cycle once, the configuration (s, v_1) is reached. Since the weight of the cycle is non-negative, $v_1 \geq v_0$. Then, by Lemma 5.4 the walk can follow the cycle one more time, reaching (s, v_2) , and still be admissible. By iterating this process, it is possible to build an infinite admissible walk.

Conversely, let $(s_0, v_0) \rightarrow \dots \rightarrow (s_n, v_n) \rightarrow \dots$ be an infinite admissible walk. Since there are only finitely many vertices, there exists at least one vertex s' appearing infinitely many times in it. Let (s'_i, v'_i) be the occurrences of the corresponding configurations in the walk. Since the component-wise order over vectors of \mathbb{N}^k is a well partial order, there exists i, j such that $v'_i \leq v'_j$. The cycle followed between s'_i and s'_j has a positive weight. \square

5.2 Decidability of the uniform termination

Definition 5.7 (Linear parts, semi-linear parts). Let $(M, +)$ be a commutative monoid. A *linear part* of M is a subset of the form $v + V^*$ where $v \in M$ and V is a finite subset of

M . That is, if $V = \{v_1, \dots, v_p\}$, a linear part can be expressed as:

$$\left\{ v + \sum_{i=1}^{i=p} n_i v_i \mid n_i \in \mathbb{N} \right\}$$

A *semi-linear part* of M is a finite union of linear parts.

LEMMA 5.8. *In a commutative monoid, rational parts are exactly the semi-linear parts.*

Recall that rational parts are build from $+$, union and Kleene's star $*$. When dealing with words (that is the free monoid generated by a finite alphabet), $+$ is word concatenation (not commutative) and so rational parts are exactly the regular languages.

PROOF. Semi-linear parts are expressed as rational parts.

Conversely, it is sufficient to show that the set of semi-linear parts contains all finite parts and is closed by union, sum and $*$. The hard point being the closure under $*$ which is a consequence of commutativity. It holds because $(v + V^*)^* = (v + (\{v\} \cup V)^*) \cup \{0\}$ (the key idea being that $(a(b^*))^* = a^*b^*$ in a commutative monoid). See [Reutenauer 1989] (Proposition 3.5) for details. \square

LEMMA 5.9. *The set of cycles in a graph is a rational part (of the free monoid generated by the edges).*

PROOF. Consider the graph as an automaton with each edge labelled by a separate label. The set of paths between two given vertices is a regular language (accepted by the automaton with the proper input and accepting nodes). So is the set of cycles as finite union of regular languages. \square

COROLLARY 5.10. *The set of weights of cycles in a VASS is a semi-linear part of \mathbb{Z}^k .*

PROOF. Since the weighting function ω is a morphism between (A, \cdot) and $(\mathbb{Z}^k, +)$, it preserves rational parts. Hence, the set of weights of cycles is a rational part of \mathbb{Z}^k . Since $+$ is commutative, it is also a semi-linear part. \square

Notice that the proofs are constructive. Hence the semi-linear part can effectively be built.

THEOREM 5.11. *Uniform termination of VASS is in NPTIME.*

PROOF. By Theorem 5.6, a VASS is *not* uniformly terminating if and only if there is a cycle whose weight is in \mathbb{N}^k . Since the set of weights of cycles is a semi-linear part of \mathbb{Z}^k , it is sufficient to be able to decide whether a linear part of \mathbb{Z}^k intersects \mathbb{N}^k (and try this for each linear part of the union).

Let $U = \{u_1, \dots, u_p\}$ and $u + U^*$ be a linear part of \mathbb{Z}^k . It intersects \mathbb{N}^k if and only if there exists $n_1, \dots, n_p \in \mathbb{N}$ such that $u + \sum n_i u_i \geq 0$.

This can be solved in NPTIME using usual integer linear programming techniques. \square

Since VASS and Petri nets are equivalent, this also shows that uniform termination of Petri nets is decidable. Without going through the equivalence, a direct and simpler proof can be made for Petri nets. Such a proof can be found in [Moyen 2003], (theorem 60, page 83).

Example 5.12. Figure 7 displays two VASS. More formally, the first one should be described as a graph $G = (S, A)$ with:




Fig. 7. Two VASS

- $S = \{a, b, c, d\}$
- $A = \{a \xrightarrow{a_1} b, a \xrightarrow{a_2} c, b \xrightarrow{a_3} c, c \xrightarrow{a_4} d, d \xrightarrow{a_5} a\}$
- $\omega(a_1) = (0, -1, +1)$, $\omega(a_2) = (-1, 0, 0)$, $\omega(a_3) = (0, +1, -2)$, $\omega(a_4) = (0, 0, +1)$, $\omega(a_5) = (0, 0, -1)$.

Let $A = (((a_1 a_3) | a_2) a_4 a_5)^*$ be a regular expression describing cycles from a to a . Then the set of all cycles in these VASS is the language recognised by the regular expression:

$$A|(a_3 a_4 a_5 A a_1)|(a_4 a_5 A(a_1 a_3 | a_2))|(a_5 A(a_1 a_3 | a_2) a_4)$$

This corresponds to the semi-linear set $\{(a_1 a_3 a_4 a_5), (a_2 a_4 a_5)\}^*$

The set of weights of cycles in the first VASS is obtained from the set of cycles by applying the weighting function (which is a morphism). Thus, we obtain:

$$\{(\omega(a_1) + \omega(a_3) + \omega(a_4) + \omega(a_5)), (\omega(a_2) + \omega(a_4) + \omega(a_5))\}^* = \{(-l, 0, -k) | k, l \in \mathbb{N}\}$$

Obviously, this set does not intersect \mathbb{N}^3 (apart from the trivial 0 solution), hence the VASS is uniformly terminating.

For the second VASS, we obtain the set $\{(3l - k, k - l, k - 2l) | k, l \in \mathbb{N}\}$. It intersects \mathbb{N}^3 , for example with $k = 2, l = 1$, corresponding to the cycle $(a_1 a_3 a_4 a_5)^2 (a_2 a_4 a_5)$ whose weight is $(1, 1, 0)$. Hence, the VASS is not uniformly terminating.

However, any infinite walk starting from the configuration $(a, (0, 14, 0))$ is not admissible. Deciding whether a given configuration leads to an infinite admissible walk or not is a different problem than uniform termination.

It is worth noticing that in the second case, the cycle detected is *not* a simple cycle. So the problem is different from the one of detecting simple cycles in graphs and requires a specific solution.

5.3 VASS as Resource Control Graphs

Before the formal definition of Resource Control Graphs, we show here how VASS can be used to build proofs of uniform termination of programs.

In the rest of this section, we consider the following size function:

$$\|\langle \text{IP}, \sigma \rangle\| = (|\text{stk}_1|, \dots, |\text{stk}_s|)_{\text{stk}_i \in \mathcal{S}}$$

that is, the vector whose components are the length of the different stacks of a given program. Moreover, we use (e_i) to denote the canonical basis of \mathbb{Z}^k , that is e_i is the vector whose j th component is $\delta_{i,j}$.




Fig. 8. The Resource Control VASS for the reverse program

Definition 5.13 (Weights). To each instruction, we assign the following *weight*:

- $\omega(\mathbf{r} := \text{pop}(\text{stk}_i)) = -e_i$
- $\omega(\text{push}(\mathbf{r}, \text{stk}_i)) = e_i$
- $\omega(i) = 0$ for all other instructions.

Definition 5.14 (Resource Control VASS). Let p be a program. Its *Resource Control VASS* is a VASS whose underlying graph is the Control Flow Graph of p and edge i has weight $\omega(i)$ as defined above.

PROPOSITION 5.15. *Let p be a program and G be its Resource Control VASS. If $\theta_0 = \langle \text{IP}_0, \sigma_0 \rangle \xrightarrow{*} \langle \text{IP}_n, \sigma_n \rangle = \theta_n$ is an execution of p , then $(\text{IP}_0, \|\theta_0\|) \xrightarrow{*} (\text{IP}_n, \|\theta_n\|)$ is an admissible walk of G with the same trace.*

PROOF. By induction on the length of the execution. Notice that executions leading to errors (\perp) are not taken into account here. \square

THEOREM 5.16. *Let p be a program and G be its Resource Control VASS. If G is uniformly terminating, then p is uniformly terminating.*

PROOF. Otherwise, there would exist an infinitely long execution that can be mapped onto an infinite admissible walk by the previous Proposition. \square

Since uniform termination of VASS is decidable, this allows to decide uniform termination of a broad class of programs. Of course, the converse is not true since uniform termination of programs is not decidable.

Example 5.17. The Resource Control VASS of the reverse program is displayed on Figure 8. Since it is uniformly terminating, so is the reverse program.

Weighted graphs, as used in Section 4 to prove Non-Size Increasingness of programs are also a special case of VASS with only one dimension.

6. RESOURCE SYSTEMS WITH STATES

Resource Systems with States (RSS) are a generalisation of VASS seen in the previous Section. In VASS, the only information kept is a vector of integers, and only addition of vectors can be performed. When modelling programs, this is not sufficient. Indeed, if one wants to closely represent the memory of a stack machine, a vector is not sufficient.

Moreover, vector addition is not powerful enough to represent usual operations such as copy of a variable ($x := y$).

Hence, we will now relax the constraints on valuations and weights and basically allow valuations to be drawn from any set and weights to be any kind of functions (between valuations). Notice that for VASS, the addition of a vector v could be represented as the function $\lambda x.x + v$.

In order to be a bit more general, we will even allow the sets of valuations to be different for each vertex. This may seem strange, but a typical use of that is to have vectors with different numbers of components as valuations (that is the set of valuations for vertex s_i would be \mathbb{Z}^{k_i}) and matrix multiplications as weights (where the matrices have the correct number of rows and columns). Of course, it is always possible to take the (disjoint) union of these sets, but it usually clutters needlessly the notations. See Example 9.3 for more details.

6.1 Graphs and States

Definition 6.1 (RSS, configurations, walks). A *Resource System with States* (RSS) is a tuple (G, V, V^+, W, ω) where

- $G = (S, A)$ is a directed graph, $S = \{s_1, \dots, s_n\}$ is the set of vertices and $A = \{a_1, \dots, a_m\}$ is the set of edges.
- V_1, \dots, V_n are the sets of *valuations*. V is the union of all of them.
- $V_i^+ \subset V_i$ are the sets of *admissible valuations*. V^+ is the union of them.
- $W_{i,j} : V_i \rightarrow V_j$ are the sets of *weights*. W is the union of them.
- $\omega : A \rightarrow W$ is the *weighting function* such that $\omega(a) \in W_{i,j}$ if $s_i \xrightarrow{a} s_j$.

When both the valuations and weights sets are clear, we will name the RSS after the underlying graph G .

A *configuration* is a couple $\eta = (s, v)$ where $s = s_i \in S$ is a vertex of the graph and $v \in V_i$ is a valuation. A configuration is *admissible* if $v \in V_i^+$ is admissible.

A *walk* is a sequence (finite or not) of configurations $(s_0, v_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, v_n) \xrightarrow{a_{n+1}} \dots$ such that $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n \xrightarrow{a_{n+1}} \dots$ and for all $i > 0$, $v_i = \omega(a_i)(v_{i-1})$. A walk is *admissible* if all configurations in it are admissible.

The walk *follows* path p which is called either *underlying path* or *trace* of the walk.

As earlier, we write $\eta \rightarrow \eta'$ if the relation holds for an unspecified edge and $\xrightarrow{+}, \xrightarrow{*}$ for the transitive and reflexive-transitive closures.

The idea behind having both valuations and admissible valuations is that this allows V to have some nice algebraic properties not shared by V^+ . Moreover, this also allows the set of valuations to be the closure of the admissible valuations under the weighting functions, thus removing the deadlock problem of reaching something that would not be a valuation (and replacing it by the more semantical problem of detecting non admissible valuations). Typically with VASS, V is \mathbb{Z}^k , thus being a ring, and V^+ is \mathbb{N}^k . Since weights can add any vector, with positive or negative components, to a valuation, V is the closure of V^+ by this operation. Moreover, VASS do not suffer from the deadlock problems that appear in Petri nets (but this is done by introducing the problem of deciding if a walk is admissible).

Notice that either unions (for V , V^+ or W) can be considered to be a disjoint union without loss of generality.

Definition 6.2 (Weight of a path). Let G be a RSS. The weighting function can be canonically extended over all paths in G by choosing $\omega(ab) = \omega(b) \circ \omega(a)$.

(W, \circ) is a magma. It is not a monoid because the identity is not unique. There is a finite set of neutral elements, the identities over each V_i .

Notice that we do not actually need the whole W . Only the part generated by the individual weights of edges is necessary to handle a RSS. We will overload the notation and call it W as well.

In practise, it is often more convenient to describe W as a set together with some right-action on V . That is, there is an operation $\otimes : V \times W \rightarrow V$ such that $v \otimes \omega(a) = \omega(a)(v)$. In this case, the function composition becomes an internal law of W , $\S : W \times W \rightarrow W$ such that $\omega(a) \S \omega(b) = \omega(b) \circ \omega(a)$. This turns ω in a morphism between (A, \cdot) and (W, \S) .

This notation is much more convenient when composing weights along a path. Indeed, since ω is a morphism, $\omega(ab) = \omega(a) \S \omega(b)$, that is the weights are composed in the same order as the edges along the path while using functional composition we had $\omega(ab) = \omega(b) \circ \omega(a)$, needing to reverse the order of edges along the path. Moreover, since weights usually have some common shape, (W, \S) is usually a well known algebraic structure.

Example 6.3. For the VASS of previous Section, we have $V_i = \mathbb{Z}^k$ and $V_i^+ = \mathbb{N}^k$ for all i , and $\omega(a_i) = \lambda x.x + u_i$ for some vector $u_i \in \mathbb{Z}^k$. Or, we could describe VASS by saying that $V = W = \mathbb{Z}^k$, $V^+ = \mathbb{N}^k$ and $\otimes = \S = +$.

The notation with \otimes and \S is much more convenient, especially to handle easily weights of paths such as done in the lemmas and theorems of the previous Section.

Moreover, the fact that weights (as functions) all have the same shape (namely, $\lambda x.x + \alpha$) allows to identify each weight with the vector α , thus giving a more convenient definition.

Of course, If we consider V_i as objects and $\omega \in W$ as arrows, we have a category. Indeed, identity exists for each V_i and composition of two arrows is properly defined.

6.2 Properties of RSS

6.2.1 Order

Definition 6.4 (Ordered RSS). An ordered RSS is an RSS $G = (G, V, V^+, W, \omega)$ together with a partial ordering \prec over valuations such that the restriction of \prec over V^+ is a well partial order.

For VASS, the component-wise order on vectors of the same length is the well partial order (over $V^+ = \mathbb{N}^k$) that was used in the previous Section.

Definition 6.5 (Monotonicity, positivity). Let (G, V, V^+, W, ω) be an ordered RSS. We say that it is *monotonic* if all weighting functions $\omega(a_i)$ are increasing with respect to \prec . Since the composition of increasing functions is still increasing, the weighting function of any path will be increasing.

We say that (G, V, V^+, W, ω) is *positive* if for each $v \in V^+$ and $v' \in V$, $v \prec v'$ implies $v' \in V^+$.

VASS are both monotonic and positive. Monotonicity is the key of Lemma 5.4 while positivity is implicitly used in the proof of Theorem 5.6 to say that the valuation reached after one cycle is still admissible.

Definition 6.6 (Resource awareness). Let G be an ordered RSS and $f : V \rightarrow V$ be a function. G is f -resource aware if for any walk $(s_0, v_0) \xrightarrow{*} (s_n, v_n)$ we have $v_n \preceq f(v_0)$

6.2.2 Uniform termination

Definition 6.7 (Uniform termination). Let G be a RSS. G is *uniformly terminating* if there is no infinite admissible walk in G .

Notice that if a RSS is not uniformly terminating, then there exists an infinite admissible walk that stay entirely within one strongly connected component of the underlying graph. In the following, when dealing with infinite walks we suppose without loss of generality that the RSS is strongly connected.

Theorem 5.6 can be generalised to RSS:

THEOREM 6.8. *If G does not uniformly terminates, then there is an admissible cycle $(s, v) \xrightarrow{+}(s, u)$ with $v \preceq u$. If G is monotonic and positive, then this is an equivalence.*

PROOF. If an infinite admissible walk exists, then we can extract from it an infinite sequence of admissible configurations (s', v_k) since there is only a finite number of vertices. Since the order is a well partial order on V^+ , there exists a $i < j$ with $v_i \preceq v_j$, thus leading to the cycle.

If the cycle exists, then it is sufficient to follow it infinitely many time to have an infinite admissible walk. Monotonicity is needed to ensure that every time one follows the cycle, the valuation does indeed increase. Positivity is needed to ensure that when going through always increasing valuations one will never leave V^+ . \square

PROPOSITION 6.9. *Let $G = (G, V, V^+, W, \omega)$ be a RSS.*

- (1) *If V is finite, then W is finite.*
- (2) *If V is finite, then uniform termination of G is decidable.*
- (3) *If both V and W are enumerable, then it is semi-decidable to know if an ordered RSS G is not uniformly terminating.*

PROOF.

- (1) Because the set of functions $\mathcal{F}(V, V)$ is finite and contains W .
- (2) If there are only finitely many valuations, the cycle of Theorem 6.8 becomes $(s, v) \xrightarrow{+}(s, v)$, that is one comes back to exactly the same configuration. Then it is possible to compute all the possible weights of cycles (there are only finitely many of them) and check with all the valuation if the condition is met. Notice that this does not require the RSS to be ordered.
- (3) By enumerating the cycles and the valuations simultaneously, computing the new valuation after going through the cycle and checking with the ordering if this satisfies Theorem 6.8.

\square

Corollary 5.10 can be generalised:

PROPOSITION 6.10. *If (W, \wp) is commutative, then the set of weights of cycles of a RSS is semi-linear.*

This will allow to easily find candidates for a generalisation of Theorem 5.11 if the set of “positive” weights is easily expressible (as it was the case for VASS). Among other, if it is itself semi-linear, then uniform termination is decidable (intersection of two semi-linear parts being decidable).

6.3 Equational versus constraint based approach

Up to now, the only weights we have considered are functions, meaning that if $s \xrightarrow{a} s'$, for each valuation v there is only one valuation v' such that $(s, v) \xrightarrow{a} (s', v')$. Sometimes, it is more convenient to have several possible results because approximations done on the values leads to a lost of knowledge. In this case, the weights considered will be relations rather than functions and we require $v' \in \widehat{\omega}(a)(v)$ rather than $v' = \omega(a)(v)$.

6.3.1 Constraints RSS

Definition 6.11 (RSS, configurations, walks).

A *Constraints RSS* is a tuple (G, V, V^+, W, ω) where

- $G = (S, A)$ is a directed graph.
- $V_i^+ \subset V_i$ are the sets of admissible valuations and valuations.
- $W_{i,j} : V_i \rightarrow \mathcal{P}(V_j)$ are the sets of *weights*.
- $\widehat{\omega} : A \rightarrow W$ is the *weighting function* such that $\widehat{\omega}(a) \in W_{i,j}$ if $s_i \xrightarrow{a} s_j$.

Configurations and admissible configurations are defined as earlier.

A *walk* is a sequence (finite or not) of configurations $(s_0, v_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, v_n) \xrightarrow{a_{n+1}} \dots$ such that $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n \xrightarrow{a_{n+1}} \dots$ and for all $i > 0$, $v_i \in \widehat{\omega}(a_i)(v_{i-1})$. A walk is *admissible* if all configurations in it are admissible.

It is important to notice that even if weighting functions return sets (that is, they are relations rather than functions), each walk has to choose one element from this set as a new valuation. That is, we do not consider configurations with sets as valuations, but rather introduce some kind of non-determinism in the RSS. The main use for this will be when some valuations are in no way related to the previous ones and can be anything (*e.g.* if a value is provided via some external mechanism such as a `scanf` instruction).

Definition 6.12 (Weight of a path). Let G be a RSS. The weighting function can be canonically extended over all paths in G by choosing $\widehat{\omega}(ab)(x) = \widehat{\omega}(b)(\widehat{\omega}(a)(x)) = \{\widehat{\omega}(b)(y) \mid y \in \widehat{\omega}(a)(x)\}$.

As earlier, uniform termination means that there exists no infinite admissible walk. However, monotonicity becomes $x \preceq y \Rightarrow \forall x' \in \widehat{\omega}(x), \exists y' \in \widehat{\omega}(y) / x' \preceq y'$.

Then, Theorem 6.8 becomes:

THEOREM 6.13. *Let G be a positive monotonic Constraints RSS. G is not uniformly terminating if and only if there is an admissible cycle $(s, v_0) \xrightarrow{+} (s, v_1)$ such that $v_0 \preceq v_1$.*

PROOF. If an admissible infinite walk exists, then we can extract from it an admissible cycle in exactly the same way as in Theorem 6.8.

Conversely, if an admissible cycle c exists, let $(s, v_0) \xrightarrow{a} (s', v'_0) \xrightarrow{*} (s, v_1)$ be the firsts and last configurations when following the cycle. By hypothesis, $v_0 \preceq v_1$.

Then, there exists $v'_1 \in \widehat{\omega}(a)(v_1)$ such that $(s, v_1) \xrightarrow{a} (s', v'_1)$ and $v'_0 \preceq v'_1$. By positivity of the VASS, v'_1 is still admissible.




Fig. 9. Constraints VASS for Ackermann's function

By iterating this process, we build the admissible cycle $(s, v_1) \xrightarrow{c} (s, v_2)$ with $v_1 \preceq v_2$. Then, this can be done *ad infinitum* thus leading to an admissible infinite walk. \square

6.3.2 Constraints VASS.

Let us show how this concept apply to VASS and why it can be useful when studying programs.

Definition 6.14 (Constraints VASS). A *Constraints VASS* is a directed graph $G = (S, A)$ together with a *weighting function* $\omega : A \rightarrow \overline{\mathbb{Z}^k}$ where k is a fixed integer.

A *configuration* is a couple $\eta = (s, v)$ where $s \in S$ and $v \in \mathbb{Z}^k$. It is *admissible* if $v \in \mathbb{N}^k$.

A *walk* is a sequence (finite or not) of configurations $(s_0, v_0) \xrightarrow{a_1} \dots \xrightarrow{a_n} (s_n, v_n)$ such that $s_0 \xrightarrow{a_1} s_1 \xrightarrow{a_2} \dots \xrightarrow{a_n} s_n$ and for all $i > 0$, $v_i \leq v_{i-1} + \omega(a_i)$. A walk is *admissible* if all configurations in it are admissible.

To express a Constraints VASS as a Constraints RSS, we should consider the weighting function $\widehat{\omega}(a) : \mathbb{Z}^k \rightarrow \mathcal{P}(\mathbb{Z}^k)$ such that $\widehat{\omega}(a)(v) = \{v' \mid v' \leq v + \omega(a)\}$. Then, the relation between valuations in a walk will be the general $v_i \in \widehat{\omega}(a_i)(v_{i-1})$. Since, all constraints have the same shape, we can express this in a more readable way. Constraints VASS are positive and monotonic. When there is no $+\infty$ in the weights, it is always “best” to choose the greatest possible valuation, that is use the (regular) VASS with the same underlying graph and weighting function.

Example 6.15. Consider the following functional program computing Ackermann's function:

$$\begin{aligned} \text{Ack}(0, n) &\rightarrow n + 1 \\ \text{Ack}(m + 1, 0) &\rightarrow \text{Ack}(m, 1) \\ \text{Ack}(m + 1, n + 1) &\rightarrow \text{Ack}(m, \text{Ack}(m + 1, n)) \end{aligned}$$

For functional programs, an equivalent of the CFG can be the *calls graph*. there is one vertex for each function symbol (here only one) and one edge for each call (here, 3). Since there are two positive integers in the program, it is natural to choose (m, n) as valuation.

However, when considering the outer call in the last line the second argument is $\text{Ack}(m + 1, n)$ which cannot be related to the parameter n in any easy way. So, using a regular VASS, this call would not be representable, while, *e.g.*, the call in the second line corresponds to adding the vector $(-1, +1)$ to the valuation.

With a Constraints VASS, we can represent this last call. Indeed, not knowing anything on the result simply means that we can relax all constraints on it which will be represented by the vector $(-1, +\infty)$. The constraints VASS for Ackermann's function is displayed on Figure 9.

Since this Constraints VASS is uniformly terminating, so is Ackermann's function.

This both illustrates why Constraints VASS can be useful as well as hints how to apply the ideas behind RCG to functional programs.

7. RESOURCES CONTROL GRAPHS

Instead of the weighted graphs or VASS used before, we will now use any RSS to model programs. A set of admissible valuations will be given to each state and weighting functions simulate the corresponding instruction.

Since we can now have any approximation of the memory (the stores) for valuations, we cannot simply use the length of a state. Instead, we consider given a *size function* that associates to each state (or to each store) some size. The size function is unspecified in general. Of course, when using RCG to model programs, the first thing to do is usually to determine a suitable size function (according to the studied property). Notice that depending on the size function, weights of instructions can or cannot be defined properly (that is, some sizes are either too restrictive or too loose and no function can accurately reproduce on the size the effect of a given instruction on actual data). In this case, the RCG cannot be defined and another size function has to be considered.

7.1 Resources Control Graphs

Definition 7.1 (RCG). Let p be a program and G be its control flow graph. Let V^+ be a set of admissible valuations (and \prec be a well partial order on it). Let $\|\bullet\| : \Theta \rightarrow V^+$ be a size function from states to valuations and V_{lb1}^+ be the image by $\|\bullet\|$ of all states $\langle \text{lb1}, \sigma \rangle$ for all stores σ .

For each i edge of G , let $\omega(i)$ be a function such that for all states θ verifying $p \vdash \theta \xrightarrow{i} \theta'$, $\omega(i)(\|\theta\|) = \|\theta'\|$. Let V be the closure of V^+ by all the weighting functions $\omega(i)$.

The *Resource Control Graph (RCG)* of p is the RSS build on G with weights $\omega(i)$ for each edge i , valuations V and admissible valuations V^+ (ordered by \prec). V_{lb1}^+ being the admissible valuations for vertex lb1 .

As stated before, we will write $v \otimes \omega(i)$ instead of $\omega(i)(v)$ and $\omega(i) \circledast \omega(j)$ instead of $\omega(j) \circ \omega(i)$.

LEMMA 7.2. *Let p be a program, G be its RCG and $p \vdash \theta_0 \rightarrow \dots \rightarrow \theta_n$ be an execution with trace t . There exists an admissible walk $(s_0, \|\theta_0\|) \rightarrow \dots \rightarrow (s_n, \|\theta_n\|)$ with the same trace t .*

THEOREM 7.3. *Let p be a program and G be its RCG. If G is uniformly terminating, then p is also uniformly terminating.*

Example 7.4. A Space-RCG as defined in Section 4 is a special case of general RCG. In this case, $\|\theta\| = |\theta|$, this leads to $V_{\text{lb1}}^+ = V^+ = \mathbb{N}$ for each label lb1 . Similarly, $\omega(i) = \lambda x.x + k_i$ with k_i as in definition 4.1. Since $k \in \mathbb{Z}$, the closure of V^+ by the weighting functions is $V = \mathbb{Z}$.

In this case, resource awareness of the Space-RCG (or β -Space-RCG) guarantees a resource bound on the program execution.

Example 7.5. For a better representation of programs, the size can be the vector where each component is the length of a stack: $\|\langle \text{IP}, \sigma \rangle\| = (|\text{stk}_1|, \dots, |\text{stk}_s|)_{\text{stk}_i \in \mathcal{S}}$. This corresponds exactly to what is done with the Resource Control VASS of Section 5.3. As shown, this allows to decide uniform termination of several programs.

This termination analysis is close to the Size Change Termination [Lee et al. 2001] in the sense that the size of data is monitored and a well ordering on it ensure that it cannot decrease forever. It is sufficient to prove uniform termination of most common lists programs such as reversing a list or insertion sort. It is also, in some way, slightly more efficient than the original SCT because it can take into account not only the decreasing in size, but also the increasing. In this way, a program that would loop on something like `pop pop push` (2 pops and 1 push) is not caught by SCT but is proved uniformly terminating with this analysis. In this sense, it is closer to the SCT with difference constraints (δ SCT) [Ben-Amram 2006].

This method is in NPTIME, as we have shown, uniform termination of VASS is in NPTIME. The original SCT, as well as fan-in free δ SCT, is PSPACE-complete. However, this simple method does not allow for data duplication or copy. Lee, Jones and Ben-Amram already claimed in the original SCT that there exists a poly-time algorithm for SCT dealing with “programs whose size-change graphs have in- and out-degrees bounded by 1”. It is easy to check that VASS can only model such kind of programs accurately⁵, hence the NP bound is not a big surprise.

Moreover, this method has a fixed definition of size and hence will not detect termination of programs whose termination argument does not depend on the decrease of the length of a list. Among other, any program working solely on integers (represented as letters of the alphabet) will not be analysed correctly.

Example 7.6. However, even this representation can be improved. Typically, using Resource Control VASS it is impossible to detect anything happening to registers. If we have a suitable size function $\|\bullet\| : \Sigma \rightarrow \mathbb{N}$ for registers⁶, we can choose $\|\langle \text{IP}, \sigma \rangle\| = (\|\mathbf{r}_1\|, \dots, \|\mathbf{r}_r\|)_{\mathbf{r}_i \in \mathcal{R}}$. In this case, depending on the operators, weight could be either vectors addition or matrices multiplication (to allow the copy of a register).

Remark 7.7. Taking exactly the image of $\|\bullet\|$ as the set of admissible valuations V^+ might be a bit too harsh. Indeed, this set might have any shape and is probably not really easy to handle. So, it is sometimes more convenient to consider a superset of it in order to easily decide if a valuation is admissible or not. The convex hull (in V) of the image of $\|\bullet\|$ is typically such a superset. Notice that it is very similar to the idea of trying to find an admissible set of sequences of states which will be more manageable than the set of executions. Here, we try to find an admissible set of valuations which is more manageable than the actual set of sizes. For more details on how to build and manage such a superset, see the work of Avery [2006].

Remark 7.8. The size function is not specified and may depend on the property one wants to study. We do not address here the problem of finding a suitable size function for a given program. As hinted, it might be a simple vector of functions over stacks and registers but it can also be a more complicated function such as a linear combination or so. Hence,

⁵And cannot even model all those programs due to the restriction on copying variables.

⁶Note that the *size* function used here is in no way related to the *length* of a state. It plays no role when computing the space usage of a state and may also be seen as an ordering over the alphabet.

with a proper size function, one is able not only to check that a given register (seen as an integer) is always positive but also that a given register is always bigger than another one. This is similar to Avery’s functional inequalities [2006].

Example 7.9. Let us consider the following program, working on integers (that is the alphabet is the set `unsigned` of 32 bits positive integers):

<pre>0 : i := 0; 1 : if i ≥ n then goto 5; 2 : i := i + 1; 3 : some instructions modifying neither i nor n</pre>		<pre>4 : if i < n then goto 2; 5 : i := i + 1; end : end;</pre>
---	--	--

This is simply a loop `for (i=0; i<n; i++)` (in a C-like syntax). If we consider a size function that simply takes the vector of the registers, that is $\|\langle \text{IP}, \sigma \rangle\| = (\mathbf{i}, \mathbf{n})$, then the loop will have weight $(+1, 0)$ and thus lead to a cycle of positive weight. However, a clever analysis of the program could detect that inside the loop we must necessarily have $n - i > 0$ and thus suggest the size $\|\langle \text{IP}, \sigma \rangle\| = \mathbf{n} - \mathbf{i}$. Using this, the loop has weight -1 and we can prove uniform termination of the program.

As stated, we do not address here the problem of finding a correct size function for a given program. This problem is undecidable in general. But invariants can often be automatically generated, usually by looking at the pre- and post-conditions of the loops.

Notice also that this inequality must hold only in the loop. Indeed, at label 5 or after, we may have $\mathbf{i} > \mathbf{n}$. Hence using this size function everywhere would cause troubles since then $\|\langle 5, \sigma \rangle\|$ will not be admissible.

Having different sets of valuations for each labels, that is a size function operating differently on each label, can solve this problem. By choosing $\|\langle \text{IP}, \sigma \rangle\| = (\mathbf{i}, \mathbf{n})$ for `IP = 0, 1, 5, end` and $\|\langle \text{IP}, \sigma \rangle\| = (\mathbf{i}, \mathbf{n}, \mathbf{n} - \mathbf{i})$ otherwise, we can ensure that the “natural” sets of admissible valuations (\mathbb{N}^2 and \mathbb{N}^3) indeed correspond to the image of the size function (or at least a manageable superset of it).

In this case, of course, we need the weight between labels 1 and 2 to take into account the apparition of a new component in the valuation. Here, this can be done using a matrix multiplication since the new component in the valuation is a linear combination of the existing ones. See Example 9.3 for the complete construction of the RCG.

7.2 Constraints RCG

Constraints RSS can also be used instead of RSS to model programs and build RCG as was done with the Ackermann’s function of Example 6.15. In that case, the relation required between weights and sizes is:

for all states θ verifying $p \vdash \theta \xrightarrow{i} \theta'$, $\|\theta'\| \in \widehat{\omega}(i)(\|\theta\|)$.

Then, the simulation Lemma and uniform termination Theorem are still true:

LEMMA 7.10. *Let p be a program, G be its Constraints RCG and $p \vdash \theta_0 \rightarrow \dots \rightarrow \theta_n$ be an execution with trace t . There exists an admissible walk $(s_0, \|\theta_0\|) \rightarrow \dots \rightarrow (s_n, \|\theta_n\|)$ with the same trace t .*

PROOF. Because $\|\theta_i\|$ belongs to $\widehat{\omega}(a)(\theta_{i-1})$ and can thus always be chosen as the new valuation. \square

THEOREM 7.11. *Let p be a program and G be its RCG. If G is uniformly terminating, then p is also uniformly terminating.*

8. δ -SIZE CHANGE TERMINATION

We consider here the $(\overline{\mathbb{Z}}, \min, +)$ semi-ring and denotes \min as \oplus and $+$ as \otimes . These operations are canonically extended to define multiplication of matrices from $\mathcal{M}(\overline{\mathbb{Z}})$.

8.1 Matrices and graphs

Definition 8.1 (Sign matrix). Let M be a matrix of $\mathcal{M}(\overline{\mathbb{Z}})$. Its *sign matrix*, \overline{M} is the matrix such that $\overline{M}_{i,j}$ is $+\infty$ (resp. $0, -1, +1$) if $M_{i,j}$ is $+\infty$ (resp. $0, < 0, > 0$).

Definition 8.2 (Sign-idempotence). Let M be a square matrix of $\mathcal{M}(\overline{\mathbb{Z}})$. It is *sign-idempotent* if $\overline{M} = \overline{M} \otimes \overline{M}$, that is M has the same sign as M^2 .

Matrix M is *strongly sign idempotent* if for all $k > 0$, $\overline{M} = \overline{M^k}$, that is M has the same sign as all its powers.

Matrix M is *strongly diagonally sign idempotent* (SDSI) if for all $k > 0$, $\overline{M}_{i,i} = \overline{M^k}_{i,i}$, that is the diagonal of M has the same sign as the diagonal of all the powers of M .

Remark 8.3. Sign idempotence and strong sign idempotence are not equivalent as shown by the following matrix (remember that we are working in the $(\overline{\mathbb{Z}}, \min, +)$ semi-ring and not in the usual $(\mathbb{Z}, +, \times)$ ring):

$$M = \begin{bmatrix} 1 & -5 & -1 \\ 6 & 1 & 2 \\ 2 & -1 & 1 \end{bmatrix} \quad M^2 = \begin{bmatrix} 1 & -4 & -3 \\ 4 & 1 & 3 \\ 3 & -3 & 1 \end{bmatrix} \quad M^3 = \begin{bmatrix} -1 & -4 & -2 \\ 5 & -1 & 3 \\ 3 & -2 & -1 \end{bmatrix}$$

Definition 8.4 (Constraint graph). Let M be a square matrix of dimension n . Its *constraint graph* is a weighted directed graph G such that:

- There are n vertices $X_i, 1 \leq i \leq n$ plus an extra vertex Y .
- If $M_{i,j} \neq +\infty$, there is an edge of weight $M_{i,j}$ between X_i and X_j .
- There is an edge of weight 0 between Y and X_i , for all i .

Definition 8.5 (l -weight). Let G be a directed weighted graph. The *l -weight* between a and b is the minimum weight of all paths of length l between a and b and $+\infty$ if there is no such path.

Let $\omega_{a,b}(l)$ be the l -weight between a and b and $\overline{\omega_{a,b}(l)}$ be its sign (in $\{-1, 0, +1, +\infty\}$ as earlier).

The coefficient $M_{i,j}^k$ is the k -weight between X_i and X_j in the constraint graph of M . Before proving the Lemma,

LEMMA 8.6. *Let M be a square matrix. There exists $k > 0$ such that M^k is strongly diagonally sign idempotent.*

Let M be a square matrix. There exists $k > 0$ such that M^k is strongly sign idempotent.

The proof will be done on the constraint graph of M by looking on the l -weights. On graphs, the Lemma is expressed as follows:

LEMMA 8.7. *Let G be a directed weighted graph. There exists k such that for all vertices s, r , the set $\{\overline{\omega_{s,r}(kn)} \mid n > 0\}$ is a singleton.*

That is, for each pair of vertices, the kn -weights between them keep the same sign.

For the sake of clarity, we denote s_i the vertices of G and use $f_{i,j}(l) = \overline{\omega_{s_i,s_j}(l)}$ for the sign of the l -weights.

So, we need to show that there exists k such that for all i, j , $\{f_{i,j}(kn) | n > 0\}$ is a singleton.

PROOF. (Ben-Amram⁷)

Permuting the quantifiers. Firstly, if for all i, j , there exists $k_{i,j}$ such that $\{f_{i,j}(k_{i,j}n) | n > 0\}$ is a singleton, then it is sufficient to choose $k = \prod k_{i,j}$ to have the property.

Dealing with the diagonal. Then, consider $k_{i,i}$. If there is a cycle of negative weight from s_i to itself, let l be its length, by choosing $k_{i,i} = l$ we can ensure that $f_{i,i}(ln) = -1$ for all $n > 0$. If there is no cycle of negative weight but one of weight 0 (resp. only cycles of weight > 0), then we can again choose $k_{i,i} = l$ (the length of such a cycle) and be sure that $f_{i,i}(ln)$ is 0 (resp. $+1$). If there are no cycles from s_i to itself, then $f_{i,i}(n)$ is $+\infty$ for all n (that is, $k_{i,i} = 1$).

Then, if there is no path between s_i and s_j , then $f_{i,j}(n)$ is $+\infty$ for all n and $k_{i,j} = 1$.

In terms of matrices, this proves the first part of Lemma 8.6, that is every matrix has a power which is SDSI.

So, without loss of generality, we can consider that if there exists a cycle with a weight of a given sign, then there exists a cycle of length 1 whose weight has the same sign. Indeed, that means only considering cycles (and paths) whose length is multiple of $\prod k_{i,i}$.

In terms of matrices, that means that we now only consider SDSI matrices.

For all pairs of vertices X_i, X_j , we consider all the triples (s, C, p) such that p is a path from X_i to X_j going through s and C is a cycle from s to itself. Since we're now only considering the SDSI case, if there exists a triple (s, C, p) then there exists another one (s, C', p) such that C' is of length 1 and the weights of C and C' have the same sign. Since we will only reason on the signs of the weights, we can without loss of generality consider only triples with cycles of length 1.

Now, depending on the signs of the weights of both C and p , there can be several cases:

(1) There exists (s, C, p) with $\omega(C) < 0$.

In this case, if $k \geq k_0$ is large enough, there exist a path of length k and negative weight that loops sufficiently many times through C to “cancel” the weight of p (because the order on \mathbb{Z} is archimedean). We can choose $k_{i,j} = k_0$.

(2) The above does not hold, but there exists (s, C, p) with $\omega(C) = 0$ and $\omega(p) < 0$.

If $k \geq k_1$ the length of p , there exist a path of negative weight and length k . Again, $k_{i,j} = k_1$.

(3) None of the above hold, but there exists (s, C, p) with $\omega(C) = \omega(p) = 0$.

The same reasoning yields paths of weight 0.

(4) None of the above hold, but there exists (s, C, p) anyway.

In that case, either $\omega(C) = 0$ and $\omega(p) > 0$, in which case all paths will have strictly positive weight, or $\omega(C) > 0$ in which case all sufficiently long paths will have strictly positive weight (because they'll have to go through C sufficiently many time to “cancel” the weight of p).

(5) None of the above holds, that is there is no (s, C, p) .

In this case, either there is no path from X_i to X_j , and the corresponding sign is $+\infty$ for all powers, or there are paths but none of them is adjacent to a cycle, and there is a maximum length of a path from X_i to X_j and all further power will have sign $+\infty$.

⁷This proof improves and shortens the earlier proof of the Lemma done by the author.

□

LEMMA 8.8. *The system $X \leq X \otimes M$ has a solution if and only if there is no strictly negative coefficient in the diagonal of M^k , for all k . In that case, it admits a non-negative solution.*

PROOF. The matrix inequation corresponds to the set of inequations $X_j \leq X_i + M_{i,j}$.

If there is no strictly negative coefficient in the diagonal of M^k , that means that the constraints graph G has no cycle of strictly negative weight. In this case, we can choose for X_i the value of the shortest path to reach it from Y . This is well defined because there is no cycle of strictly negative weight and provides a solution for the system because $X_j \leq X_i + M_{i,j}$ by definition of shortest paths.

Conversely, if there is a path of strictly negative weight, then it is easy to see that by adding the inequations corresponding to the edges in this path one will eventually reach an inequation $X_i < X_i$ and the system has no solution.

If there is a solution, then $X + (1, \dots, 1)$ is also a solution. Hence, there exists a solution where all values are positive. □

8.2 Size Change Termination

We explain here how to build RCG in order to perform the same kind of analysis as the Size-Change Termination with difference constraints (δ SCT) of Ben-Amram [2006]. Here, we use matrices rather than Size Change Graphs thus following the work of Abel and Altenkirch [2002] where similar SCT matrices are used (but over a 3-valued set, thus mimicking the initial SCT and not the work with difference constraints).

In this whole section, we consider a fixed program p , and for each label $1b1_a$ in it a fixed integer k_a . Let $V_a = \mathbb{Z}^{k_a}$ and $V_a^+ = \mathbb{N}^{k_a}$ be sets of (admissible) valuations associated with each label and we consider given a size function $\|\bullet\|$ such that for each label $1b1_a$ and for each store σ , $\|\langle 1b1_a, \sigma \rangle\| \in V_a^+$.

Definition 8.9 (Size Change Matrix). Let i be an instruction in p corresponding to an edge between $1b1_a$ and $1b1_b$ in G . The *Size Change Matrix* (SCT matrix) of i is a matrix $M^{(i)}$ of $\mathcal{M}_{k_a, k_b}(\overline{\mathbb{Z}})$ such that for all states θ_a with $p \vdash \theta_a \xrightarrow{i} \theta_b$, $\|\theta_b\| \leq \|\theta_a\| \otimes M^{(i)}$.

This means that if $\|\theta_a\| = (x_1, \dots, x_{k_a})$ and $\|\theta_b\| = (y_1, \dots, y_{k_b})$, we have for each j : $y_j \leq \min_k \{x_k + M_{k,j}^{(i)}\}$ where the coefficients of $M^{(i)}$ can be any integer or $+\infty$.

Definition 8.10 (Size Change RCG). The *Size Change RCG* (SCT-RCG) of p is the Constraints RCG for p build with admissible valuations \mathbb{N}^{k_a} , and valuations \mathbb{Z}^{k_a} for vertex $1b1_a$. The weight for edge i is such that $\widehat{\omega}(i)(v) = \{v' \mid v' \leq v \otimes M^{(i)}\}$ where $M^{(i)}$ is the SCT matrix for i .

As for Constraints VASS, the common shape of constraints allows to use a weighting function $\omega(i) = M^{(i)}$ instead of the weighting relation $\widehat{\omega}$ and ask along a walk that $v_i \leq \omega(a_i)(v_{i-1})$ rather than $v_i \in \widehat{\omega}(a_i)(v_{i-1})$.

The uniform termination Theorem for Constraints RCG (Theorem 7.11) tells us that if the SCT-RCG is uniformly terminating then so is p .

SCT-RCG are both monotonic and positive, so it will be possible to apply Theorem 6.13.

THEOREM 8.11. *Let G be the SCT-RCG of p . It is uniformly terminating if and only if for all cycles c , if the corresponding matrix $M^{(c)}$ is strongly sign idempotent, then it has a*




Fig. 10. A Size Change Termination RCG.

strictly negative coefficient on the diagonal. That is:

$$G \text{ uniformly terminating} \Leftrightarrow \forall \text{ cycle } c, (M^{(c)} \text{ strongly sign idempotent} \Rightarrow \exists i/M_{i,i}^{(c)} < 0)$$

PROOF. First, consider that there exists a cycle c such that its weight M is strongly sign idempotent and all coefficients on the diagonal are ≥ 0 . Since M is strongly sign idempotent, non of its power has a strictly negative coefficient on the diagonal. Then, by Lemma 8.8, the system $X \geq X \otimes M$ admits an admissible solution. Hence, there exists an admissible cycle $(s, X) \xrightarrow{c} (s, X \otimes M)$ and by Theorem 6.13, the SCT-RCG is not uniformly terminating.

Conversely, suppose that the SCT-RCG is not uniformly terminating. In this case, by Theorem 6.13, there exists a cycle of weight M is such that $X \leq X \otimes M$ has a solution. Hence, by Lemma 8.8, no power of M has a strictly negative coefficient on the diagonal. However, by Lemma 8.6, there exists k such that M^k is strongly sign idempotent. So M^k is strongly sign idempotent but has no strictly negative coefficient on the diagonal. \square

This condition is undecidable in general. However, if the matrices are *fan-in free*, that is in each column of each SCT matrix, there is at most one non- $+\infty$ coefficient, then the problem is PSPACE-complete. See [Ben-Amram 2006] for details. Notice that in this paper, Ben-Amram uses mostly SCT graphs and not SCT matrices. The translation from one to the other is, however, quite obvious. Similarly we present here directly a condition on the cycles of the SCT-RCG without introducing the multipaths. This is close to the “graph algorithm” introduced in [Lee et al. 2001].

The simple Size Change Principle of Lee et al. [2001] can be seen as an approximation of the δ SCT principle where only labels in $\{-1, 0, +\infty\}$ are used. Since this only gives way to finitely many different SCT matrices, this is decidable in general (PSPACE-complete).

Example 8.12. Consider the following program (adapted from [Lee et al. 2001] fifth example):

0 : if $\mathbf{y} = 0$ then goto end;	5 : goto 0;
1 : if $\mathbf{x} = 0$ then goto 6;	6 : $\mathbf{x} := \mathbf{y}$;
2 : $\mathbf{a} := \mathbf{x}$;	7 : $\mathbf{y} := \mathbf{y} - 1$;
3 : $\mathbf{x} := \mathbf{y}$;	8 : goto 0;
4 : $\mathbf{y} := \mathbf{a} - 1$;	end : end;

It can be proved terminating by choosing the size function $||\theta|| = (\mathbf{x}, \mathbf{y}, \mathbf{a})$. With this size, its SCT-RCG is displayed on Figure 10. For convenience reasons, instructions 2 – 4, as well as 6 – 7 have been represented as a single edge (with a single matrix). This allows to completely forget register \mathbf{a} and so use (\mathbf{x}, \mathbf{y}) as size. Similarly, the other SCT matrices are not depicted since they are the identity matrix. Since the SCT-RCG is uniformly terminating, so is the program.

9. MORE ON MATRICES

9.1 Matrices Multiplication System with States

If we use vectors as valuations and (usual) matrices multiplication as weights, we can define Matrices Multiplication Systems with States (MMSS) in a way similar to VASS. Admissible valuations will still be the ones in \mathbb{N}^k but k is not fixed for the RSS and may depend on the current vertex.

Definition 9.1 (Matrices Multiplication System with States). A Matrices Multiplication System with States (MMSS) is a RSS $G = (G, V, V^+, W, \omega)$ where:

- $V_i = \mathbb{Z}^{k_i}$, $V_i^+ = \mathbb{N}^{k_i}$ for some constant k_i (depending on the vertex s_i).
- Weights are matrices with integer coefficients.
- $\circledast = \otimes = \times$.

Using this, it is quite easy to model copy instructions of counters machines ($\mathbf{x} := \mathbf{y}$) simply by using the correct permutation matrix as a weight. To represent increment or decrement of a counter, an operation which was quite natural with VASS, we now need a small trick known as *homogeneous coordinates*⁸. Simply represent the n counters as a $n+1$ components vector whose first component is always 1. Then, increment or decrement of a variable just becomes a linear combination of components of the vector which can perfectly be done with matrices multiplication. For example, here is how one can model the copy ($\mathbf{x} := \mathbf{y}$) and the increment ($\mathbf{x} := \mathbf{x} + 1$).

$$(1, x, y) \times \begin{bmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 1 \end{bmatrix} = (1, y, y) \quad (1, x, y) \times \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} = (1, x+1, y)$$

Example 9.2. Using homogeneous coordinates, the program of Example 8.12 has the MMSS depicted on Figure 11. Here, matrices multiplication is done on the usual $(\mathbb{Z}, +, \times)$ ring and not on the $(\overline{\mathbb{Z}}, \min, +)$ semi-ring as for SCT-RCG.

Example 9.3. Similarly, use of homogeneous coordinates allows to build a MMSS to prove uniform termination of the program of Example 7.9. It is depicted on the left part of Figure 12 (where label 3 has been omitted). The interesting thing here is the use of vectors of different lengths at different labels, thus allowing to add the constraint $\mathbf{n} - \mathbf{i} \geq 0$ only inside the loop. This example shows both the use of disjoint sets of valuations and how to work with the functional inequalities of Avery [2006].

⁸Homogeneous coordinates were originally introduced by A. F. Möbius. There are used, among other, in computer graphics for exactly the same purposes as we do here, that is representing a translation by means of matrix multiplication.




Fig. 11. MMSS as a RCG.

But there is even more. VASS are able to forbid a $x \neq 0$ branch of a test being taken in an admissible walk if x is 0 simply by decrementing x and then incrementing it immediately after. The net effect is null but if x is 0, the intermediate valuation is not admissible. This can still be done with MMSS. VASS, like Petri nets, are however not able to test if a component is empty, that is forbid the $x = 0$ branch of a test to be taken if x is not 0.

With MMSS, we can perform this test to 0. It is indeed sufficient to multiply the correct component of the valuation by -1 . If it was different from 0, then the resulting valuation will not be admissible.

So, using these tricks it is possible to perfectly model a counters machine by a MMSS: each execution of the machine will correspond to exactly one admissible walk in the MMSS and each admissible walk in the MMSS will correspond to exactly one execution of the machine.

This leads to the following theorem:

THEOREM 9.4. *Uniform termination of MMSS is not decidable.*

Example 9.5. Consider the following program, performing addition in unary (that is, repeatedly decrementing \mathbf{x} and incrementing \mathbf{y} until \mathbf{x} is 0).

0 : if $\mathbf{x} = 0$ then goto end;	3 : goto 0;
1 : $\mathbf{x} := \mathbf{x} - 1$;	end : end;
2 : $\mathbf{y} := \mathbf{y} + 1$;	

Right side of Figure 12 depicts a MMSS for this program such that there is a one-to-one correspondence between executions of the program and admissible walks of the MMSS. The size used is $(1, \mathbf{x}, \mathbf{y})$, the 1 being here because of homogeneous coordinates. Notice that we need to add an intermediate label for the $\mathbf{x} \neq 0$ branch of the test in order to generate the temporary valuation containing $\mathbf{x} - 1$, only used to force admissible walks with $\mathbf{x} = 0$ to take the other branch.

Since such a construction can be done for any counter machine (the unary addition program uses all possible instructions for counter machines) and since counter machines are Turing-complete, this shows why uniform termination of MMSS is not decidable in general.

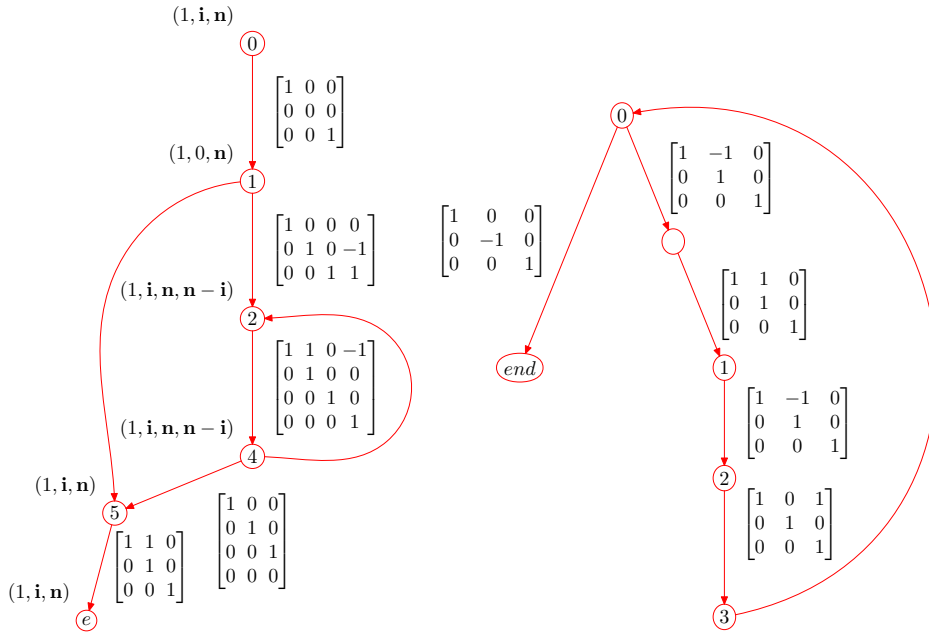


Fig. 12. MMSS for loop and unary addition.

This simulation of programs by matrices multiplications rises a surprising question. Indeed, matrices multiplications are only able to perform linear operations on data. While obviously some programs can perform non-linear operations.

This apparent contradiction is solved when we think more closely on how RSS work. Each walk in a MMSS corresponds to a matrix multiplication (because ω is a morphism), hence to a linear transformation on data. However, two different walks give rise to two different matrices, hence two different linear transformations.

When simulating a program, each different data will go through a different (admissible) walk in the MMSS. Hence, each different value will pass through a different linear transformation. Of course, the other walks (that is, the other linear transformations) also exist and are considered on this data when looking at the set of walks, but non-admissibility allows to dismiss them and only keep one.

So, from a transformation point of view, we can look at MMSS as a set of linear transformations and the admissibility mechanism selects the proper transformation to apply on each piece of data.

For example, if we consider a program performing multiplication of two integers x and y , it will likely be a loop on x , adding y to the result each time. The corresponding MMSS will have several paths (infinitely many) that can each be candidate for a walk once actual data is provided. Different paths correspond to following the loop $1, 2, 3, \dots, k, \dots$ times. Then, the walk corresponding to each of these paths will perform the linear transformation

$(1, x, y) \mapsto (1, x - k, ky)$ representable by the matrix:

$$\begin{bmatrix} 1 & -k & 0 \\ 0 & 1 & 0 \\ 0 & 0 & k \end{bmatrix}$$

However, when performing all these transformations on actual data, only those with $k \leq x$ have an admissible result and only the one with $k = x$ has all its intermediate valuations admissible. So, the admissibility mechanism selects the right linear transformation to apply.

That means that when simulating a program computing a (non-linear) function by a MMSS, the simulation actually consider the function as being *piecewise linear*, computes the result of all the possible linear transformations implied and selects the one corresponding to current data. In general, it is possible that each linear transformation is only valid for a single value.

9.2 Tensors

Moreover, the study can go further. Indeed, using matrices of matrices (that is, tensors) we can represent the adjacency graph of a MMSS (a matrix where component (i, j) is the coefficient of the edge between vertices i and j). That is, a first order program can be represented as such kind of tensors. However, it would then be possible to uses these tensors (and tensors multiplication) in order to study second-order programs. In turn, the second order programs would probably be representable by a tensor (with more dimensions) and so one.

This would lead to a tensor algebra representing high order programs.

Example 9.6. Here is a tensor representing the MMSS of the unary addition. This is simply the connectivity matrix of the graph where each edge is itself weighted by a matrix.

$$\begin{bmatrix} 0 & \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & 0 & 0 & 0 & \begin{bmatrix} 1 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \\ 0 & 0 & \begin{bmatrix} 1 & 1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & 0 & 0 & 0 \\ 0 & 0 & 0 & \begin{bmatrix} 1 & -1 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & 0 & 0 \\ 0 & 0 & 0 & 0 & \begin{bmatrix} 1 & 0 & 1 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & 0 \\ \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}$$

9.3 Polynomial time

Another interesting approach of program analysis using matrices is the one done by Niggl and Wunderlich [2006] and Kristiansen and Jones [2005]. The programs they study are

similar to our stack machines except that the (conditional) jump is replaced by a fixed iteration structure (`loop`) where the number of iterations is bounded by the length of a given stack. It is quite easy to see that both models are very similar and can simulate one another without major trouble.

Then, they assign to each basic instruction a matrix, called a *certificate* which contains information on how to polynomially bound the size of the registers (or stacks) after the instruction by their size before executing the instruction. It appears that when sequencing instructions, the certificate for the sequence turns out to be the product of the certificates for each instruction. Certificates for loops are some kind of multiplicative closure of the certificate for the body and certificate for `if` statements are the least upper bound of the two branches.

Building the certificate of a program thus leads to a polynomial bound on the result depending on the inputs which can then be turned into a polynomial bound on the running time (depending on the shape of the loops).

So, these certificates can very well be expressed in a MMSS where the valuation would give information on the size of registers (depending on the size of the inputs of the program) and the weights of instructions will be these certificates. This will exactly be a Resources Control Graph for the program. If the program is certified, then this RCG will be polynomially resource aware.

10. CONCLUSION

We have introduced a new generic framework for studying programs. This framework is highly adaptable via the size function and can thus study several properties of programs with the same global tool. Analyses apparently quite different such as the study of Non Size Increasing programs or the Size Change Termination can quite naturally be expressed in terms of Resource Control Graphs, thus showing the adaptability of the tool.

Moreover, other analyses look like they can also be expressed in this way, thus giving hopes for a truly generic tool to express and study programs properties such as termination or complexity. It is even likely that high order could be studied that way, thus giving insights for a better comprehension of high order complexity.

Theory of algorithms is not well established. This work is really on the study of programs and not of functions. Further works in this direction will shed some light on the very nature of algorithms and hopefully give one day rise to a theoretical framework as solid as our knowledge of functions. Here, the study of MMSS and the tensors multiplication hints that a tensors algebra might be used as a mathematical background for a theory of algorithms and must then be pursued.

Acknowledgements

Many thanks to A. Ben-Amram for pointing out critical flaws in an earlier version of the proof of Theorem 5.11. Thanks also to M. Hofmann for pointing out the name “homogeneous coordinates” and its use in computer graphics.

REFERENCES

- ABEL, A. AND ALTENKIRCH, T. 2002. A Predicative Analysis of Structural Recursion. *Journal of Functional Programming* 12, 1 (Jan.), 1–41.
- AMADIO, R., COUPET-GRIMAL, S., ZILIO, S. D., AND JAKUBIEC, L. 2004. A functional scenario for bytecode
ACM Transactions on Computational Logic, Vol. V, No. N, 20YY.

- verification of resource bounds. In *Computer Science Logic, 12th International Workshop, CSL'04*. Springer, 265–279.
- ASPINALL, D. AND COMPAGNONI, A. 2003. Heap Bounded Assembly Language. *Journal of Automated Reasoning (Special Issue on Proof-Carrying Code)* 31, 261–302.
- AVERY, J. 2006. Size-change termination and bound analysis. In *Functional and Logic Programming: 8th International Symposium, FLOPS 2006*, M. Hagiya and P. Wadler, Eds. Lecture Notes in Computer Science, vol. 3945. Springer.
- BELLANTONI, S. AND COOK, S. 1992. A new recursion-theoretic characterization of the poly-time functions. *Computational Complexity* 2, 97–110.
- BEN-AMRAM, A. 2006. Size-Change Termination with Difference Constraints. *ACM Transactions on Programming Languages and Systems*. To appear.
- BONFANTE, G., MARION, J.-Y., AND MOYEN, J.-Y. 2007. Quasi-interpretation: a way to control resources. *Theoretical Computer Science*. To appear, accessible <http://www.loria.fr/~marionjy/Research/Publications/Articles/TCS.pdf>.
- COBHAM, A. 1962. The intrinsic computational difficulty of functions. In *Proceedings of the International Conference on Logic, Methodology, and Philosophy of Science*, Y. Bar-Hillel, Ed. North-Holland, Amsterdam, 24–30.
- COLSON, L. 1998. Functions versus Algorithms. *EATCS Bulletin* 65, 98–117. The logic in computer science column.
- CORMEN, T., LEISERSON, C., AND RIVEST, R. 1990. *Introduction to Algorithms*. MIT Press.
- GIRARD, J.-Y. 1987. Linear logic. *Theoretical Computer Science* 50, 1–102.
- HOFMANN, M. 1999. Linear types and Non-Size Increasing polynomial time computation. In *Proceedings of the Fourteenth IEEE Symposium on Logic in Computer Science (LICS'99)*, 464–473.
- JONES, N. 2000. The expressive power of higher order types or, life without cons. *Journal of Functional Programming* 11, 1, 55–94.
- KRISTIANSEN, L. AND JONES, N. D. 2005. The flow of data and the complexity of algorithms. In *CiE'05: New Computational Paradigms*, Cooper, Lwe, and Torenlvliet, Eds. Lecture Notes in Computer Science, vol. 3526. Springer, 263–274.
- LEE, C. S., JONES, N. D., AND BEN-AMRAM, A. M. 2001. The Size-Change Principle for Program Termination. In *Symposium on Principles of Programming Languages*. Vol. 28. ACM press, 81–92.
- LEIVANT, D. AND MARION, J.-Y. 1993. Lambda Calculus Characterizations of Poly-Time. *Fundamenta Informaticae* 19, 1,2 (Sept.), 167–184.
- MOYEN, J.-Y. 2003. Analyse de la complexité et transformation de programmes. Ph.D. thesis, University of Nancy 2.
- NIELSON, F., NIELSON, H. R., AND HANKIN, C. 1999. *Principles of Program Analysis*. Springer.
- NIGGL, K.-H. AND WUNDERLICH, H. 2006. Certifying polynomial time and linear/polynomial space for imperative programs. *SIAM Journal on Computing* 35, 5 (Mar.), 1122–1147. published electronically.
- REUTENAUER, C. 1989. *Aspects mathématiques des réseaux de Petri*. Masson.
- SHEPHERDSON, J. AND STURGIS, H. 1963. Computability of recursive functions. *Journal of the ACM* 10, 2, 217–255.

Received September 2006; revised May 2006; accepted hopefully someday