

TIGHT POLYNOMIAL BOUNDS FOR LOOP PROGRAMS IN POLYNOMIAL SPACE

AMIR M. BEN-AMRAM^a AND GEOFF HAMILTON^b

^a Department of Computer Science, Holon Institute of Technology, Israel
e-mail address: benamram.amir@gmail.com

^b School of Computing, Dublin City University, Ireland
e-mail address: geoffrey.hamilton@dcu.ie

ABSTRACT. We consider the following problem: given a program, find tight asymptotic bounds on the values of some variables at the end of the computation (or at any given program point) in terms of its input values. We focus on the case of polynomially-bounded variables, and on a weak programming language for which we have recently shown that tight bounds for polynomially-bounded variables are computable. These bounds are sets of multivariate polynomials. While their computability has been settled, the complexity of this program-analysis problem remained open. In this paper, we show the problem to be PSPACE-complete. The main contribution is a new, space-efficient analysis algorithm. This algorithm is obtained in a few steps. First, we develop an algorithm for univariate bounds, a sub-problem which is already PSPACE-hard. Then, a decision procedure for multivariate bounds is achieved by reducing this problem to the univariate case; this reduction is orthogonal to the solution of the univariate problem and uses observations on the geometry of a set of vectors that represent multivariate bounds. Finally, we transform the univariate-bound algorithm to produce multivariate bounds.

1. INTRODUCTION

A static analysis algorithm takes, as input, program code and answers a question about its possible behaviours. A standard example is to find the set of values that a variable can assume (at a given program point). The algorithm seeks a set of a particular kind (for example, intervals $[a, b]$) and is *complete* if it can always find the tightest such result (e.g., the smallest interval that contains all reachable values). A more complex analysis may establish a *relation* among the values of several variables. While such results cannot be fully computable for programs in general, there are algorithms that guarantee a sound and complete solution for a class of *weak* programs, defined by removing or weakening some features of full programming languages. For example, [HOP⁺18] shows an algorithm that, for a particular weak language, finds any *algebraic relation* that exists among variable values. Such an algorithm can establish facts like: two variables x and y always satisfy, at the conclusion of the program, the equation $y = x^2$. By retaining a copy of the input, this allows

Key words and phrases: asymptotically-tight, multivariate, worst-case, polynomial bounds.

us to answer a question like: does the program square its input, or in general, can an output value be expressed as a polynomial in the input values and what is this polynomial?

In this paper we are interested in final values that are not necessarily expressible by polynomials, but are *polynomially bounded*, and we are seeking a *tight asymptotic upper bound*. A prototypical context where such a question is asked is *complexity analysis*: the quantity to be analyzed may represent a counter of steps performed by the program and we wish to establish a result such as “the program’s worst case is $\Theta(n^2)$ steps” (where n is the size of an input parameter).

In 2008, Ben-Amram et al. [BJK08] considered a simple imperative programming language with explicitly-bounded loops, and proved that it is decidable whether computed values are polynomially bounded. This work left open the problem of finding an explicit *tight* bound for such values. Recently, in [BAH20], we solved this *bound analysis problem* (the class of programs remained as in [BJK08]). However, the solution was very inefficient (doubly-exponential space) and did not settle the complexity of this problem. In this paper we answer this question and prove that tight polynomial bounds (for the same language) can be computed in polynomial space. Our approach is to first develop an analysis for *univariate* bounds only. This simplifies the task and makes the problem of solving it efficiently more tractable. However the full problem is to compute multivariate bounds (such as $O(xy + yz^2)$, equivalently $O(\max(xy, yz^2))$). In Sections 11 and 12 we show how a solution to the multivariate problem can be evolved from our solution to the univariate problem.

We now discuss, somewhat informally, some key features of the approach. We have already noted that bounds may be multivariate, i.e., depend on multiple variables. In [BAH20] we argued that it is also necessary to compute *simultaneous bounds* on multiple variables; moreover, we find sets of simultaneous bounds. The capability of computing sets of simultaneous bounds is important even if we are ultimately interested in bounding the value of a single variable. Consider, for example, the following piece of code, where `choose` represents non-deterministic choice:

```
choose Y := X*X or { Y := X; X := X*X }
```

For it we have two incomparable simultaneous bounds on X, Y : $\langle x, x^2 \rangle$ and $\langle x^2, x \rangle$ (where x stands for the initial value of X). This allows us to deduce that if the above command is followed by:

```
X := X*Y
```

The final value of X is $O(x^3)$. If, in the previous phase, we had only computed separate worst-case bounds, they would have been x^2 for each of the variables, and led to a loose bound of x^4 .

A remarkable corollary of the analysis in [BAH20] is that for our class of programs, when there is a polynomial (simultaneous) upper bound on a set of variables, it can always be matched by lower bounds; in other words, a *tight* worst-case bound can always be found. This means that the problem of computing asymptotic upper bounds and that of computing worst-case lower bounds are the same problem. For technical reasons, we formulate our results in terms of worst-case lower bounds. More precisely, we seek functions such that, for almost all inputs, for some execution, the output reaches (or surpasses) this function (up to a constant factor). We call such a function *attainable* (for a full, precise definition see Section 3). Thus, the program property that our analyzer computes is *the set of all attainable polynomial bounds for all polynomially-bounded program variables*. These bounds

are expressed by *abstract polynomials*, which are ordinary polynomials stripped of their coefficients. This agrees, of course, with the intention of providing asymptotic (“big-Oh” or “big-Omega”) bounds, where coefficients do not matter. It is not hard to see that due to this abstraction, the set of attainable bounds (for polynomially-bounded variables) becomes finite.

Even if finite, the size of this set, and the expression of its members, may be a concern; indeed, in [BAH20] we argued that the size as well as the number of maximal attainable bounds may be exponential in the program size. Here, we overcome this problem by showing that it is always possible to do with expressions of polynomial bit-size, even if the *size of the set of attainable bounds* may be exponential. This is achieved by redefining the task of our analyzer to either *verify* a given bound, or *generate* bounds—one at a time. We also design the algorithm in a new way for achieving the desired complexity. In fact, at first we could only see how to achieve such a result for univariate bounds. We later saw that it is possible to extend the solution to multivariate bounds, by exploring a reduction of the multivariate-bound problem to the univariate-bound problem. We present our algorithms in the same order, so that the justification for each step may be understood.

2. PRELIMINARIES

In this section we make some technical definitions and present the programming language that we study.

2.1. Notations. The set $[n]$ is $\{1, \dots, n\}$. For a set S an n -tuple over S is a mapping from $[n]$ to S ; the set of n -tuples is denoted by S^n . Natural liftings of operators to collections of objects are used, e.g., if t is a tuple of integers then $t + 1$ is formed by adding 1 to each component. The intention should usually be clear from context. If S is ordered, we extend the ordering to S^n by comparing tuples element-wise (this leads to a partial order, in general, e.g., with natural numbers, $\langle 1, 3 \rangle$ and $\langle 2, 2 \rangle$ are incomparable).

A binary relation on a set S is a subset $R \subseteq S \times S$ and we use the customary notation xRy instead of $(x, y) \in R$. Functions are a subclass of relations, and we liberally switch between functional notation $y = f(x)$ and relational notation.

The post-fix substitution operator $[a/v]$ may be applied to any sort of expression containing a variable v , to substitute a instead; e.g., $(x^2 + yx + y)[2z/y] = x^2 + 2zx + 2z$.

When we write “polynomials,” we are referring, unless stated otherwise, to multivariate polynomials in x_1, \dots, x_n that have non-negative integer coefficients. We presume that an implementation of the algorithm represents a polynomial as a set of monomials (which may be empty—for the zero polynomial), and that numbers are stored in binary, but otherwise we do not concern ourselves with low-level implementation details. According to context, a polynomial may refer to the function it expresses rather than the syntactic entity.

2.2. The core language. Following previous publications [BJK08, BA10, BAH20], we study the bound analysis problem for programs in a weak programming language, a “core language.” This is an imperative language, including bounded loops, non-deterministic branches and restricted arithmetic expressions; the syntax is shown in Figure 1.

$$\begin{array}{lcl}
\mathbf{X}, \mathbf{Y} \in \text{Variable} & ::= & \mathbf{X}_1 \mid \mathbf{X}_2 \mid \mathbf{X}_3 \mid \dots \mid \mathbf{X}_n \\
\mathbf{E} \in \text{Expression} & ::= & \mathbf{X} \mid \mathbf{X} + \mathbf{Y} \mid \mathbf{X} * \mathbf{Y} \\
\mathbf{C} \in \text{Command} & ::= & \text{skip} \mid \mathbf{X} := \mathbf{E} \mid \mathbf{C}_1 ; \mathbf{C}_2 \mid \text{loop } \mathbf{X} \{ \mathbf{C} \} \mid \text{choose } \mathbf{C}_1 \text{ or } \mathbf{C}_2
\end{array}$$

Figure 1: Syntax of the core language (from Ben-Amram, Kristiansen and Jones, 2008).

A program has a finite set of integer variables, $\mathbf{X}_1, \dots, \mathbf{X}_n$; by convention, n will denote the number of variables throughout this article. In code fragments we may use other identifiers ($\mathbf{X}, \mathbf{Y}, \dots$) for convenience. Variables store non-negative integers¹.

The core language is inherently non-deterministic, a property motivated by the origin of its language as a *weakening* of a fuller imperative language. This type of weakening is related to the approach of *conservative abstraction* often employed in program analysis: the **choose** command represents a non-deterministic choice, and replaces deterministic conditionals found in real programs. This reflects the idea that the analysis conservatively abstracts any conditional command by simply ignoring the condition. Weakening the language is also necessary if we wish our analysis problem to be solvable; otherwise we would clearly be blocked by undecidability. The command **loop** \mathbf{X} $\{ \mathbf{C} \}$ repeats \mathbf{C} a (non-deterministic) number of times bounded by the value of the variable \mathbf{X} upon entry to the loop. Thus, as a conservative abstraction, it may be used to model different forms of loops (for-loops, while-loops) as long as a bound on the number of iterations, as a function of the program state on loop initiation, can be determined and expressed in the language. For convenience, we stipulate that the loop-bound variable \mathbf{X} may not be assigned a new value in the loop body (this is easily seen to be non-restrictive, since we could add a dedicated variable for the loop bound).

An *unstructured* language in which loops are implemented by conditional branching and not by structured commands may also be amenable to transformation into a form that our algorithms can handle, using some front-end analysis; we will not elaborate on this but refer the reader to [BAP16].

2.3. Semantics, more formally. We define a *program state* to be an n -tuple of non-negative integers, $\mathbf{x} = \langle x_1, \dots, x_n \rangle$, representing the values of the variables $\mathbf{X}_1, \dots, \mathbf{X}_n$. Note the different fonts used for program variables, vectors and scalars.

The semantics associates with every command \mathbf{C} over variables $\mathbf{X}_1, \dots, \mathbf{X}_n$ a relation $\llbracket \mathbf{C} \rrbracket \subseteq \mathbb{N}^n \times \mathbb{N}^n$. In the expression $\mathbf{x} \llbracket \mathbf{C} \rrbracket \mathbf{x}'$, vector \mathbf{x} (respectively \mathbf{x}') is the state before (after) the execution of \mathbf{C} .

The semantics of **skip** is the identity. The semantics of an assignment $\mathbf{X}_i := \mathbf{E}$ associates to each store \mathbf{x} a new store \mathbf{x}' obtained by replacing the component x_i by the value of the expression \mathbf{E} when evaluated over state \mathbf{x} , namely with x_j , $x_j + x_k$ or $x_j x_k$, corresponding to expressions \mathbf{X}_j , $\mathbf{X}_j + \mathbf{X}_k$ and $\mathbf{X}_j * \mathbf{X}_k$. Composite commands are described by the equations:

$$\begin{aligned}
\llbracket \mathbf{C}_1 ; \mathbf{C}_2 \rrbracket &= \llbracket \mathbf{C}_2 \rrbracket \circ \llbracket \mathbf{C}_1 \rrbracket \\
\llbracket \text{choose } \mathbf{C}_1 \text{ or } \mathbf{C}_2 \rrbracket &= \llbracket \mathbf{C}_1 \rrbracket \cup \llbracket \mathbf{C}_2 \rrbracket \\
\llbracket \text{loop } \mathbf{X}_i \{ \mathbf{C} \} \rrbracket &= \{ (\mathbf{x}, \mathbf{y}) \mid \exists i \leq x_i : \mathbf{x} \llbracket \mathbf{C} \rrbracket^i \mathbf{y} \}
\end{aligned}$$

¹An alternative definition uses signed integers, see [BAH20, remark on p. 7]; this makes little difference because of the restricted arithmetics in the language.

where $\llbracket \mathbf{C} \rrbracket^i$ represents $\llbracket \mathbf{C} \rrbracket \circ \dots \circ \llbracket \mathbf{C} \rrbracket$ (with i occurrences of $\llbracket \mathbf{C} \rrbracket$); and $\llbracket \mathbf{C} \rrbracket^0 = Id$.

2.4. Multi-polynomials. In this section we define some concepts and notations used for expressing the analysis results (including intermediate results, internal to the algorithm).

Definition 2.1. A *multi-polynomial* $\mathbf{p} = \langle \mathbf{p}[1], \dots, \mathbf{p}[n] \rangle$ is an n -tuple of polynomials (with non-negative coefficients); it represents a mapping of \mathbb{N}^n to \mathbb{N}^n . We denote by MPo1 the set of multi-polynomials, where the number of variables n is fixed by context.

A multi-polynomial \mathbf{p} expresses simultaneous bounds on the variables $X_1 \dots X_n$ in terms of their respective initial values $x_1 \dots x_n$, where each component $\mathbf{p}[i]$ represents the bound on variable X_i . For example, the command $X_4 := X_2 + X_3$ is represented by the multi-polynomial $\mathbf{p} = \langle x_1, x_2, x_3, x_2 + x_3 \rangle$.

Composition of multi-polynomials, $\mathbf{p} \circ \mathbf{q}$, is naturally defined since \mathbf{p} supplies n values for the n variables of \mathbf{q} , and we have: $(\mathbf{p} \circ \mathbf{q})[i] = \mathbf{p}[i] \circ \mathbf{q}$. We define Id to be the identity transformation, $\mathbf{x}' = \mathbf{x}$ (in MP notation: $\mathbf{p}[i] = x_i$ for $i = 1, \dots, n$).

Definition 2.2 (Univariate MP). A univariate MP is defined as an n -tuple \mathbf{p} where each component $\mathbf{p}[i]$ is a polynomial in the single variable x .

Definition 2.3. APo1 , the set of *abstract polynomials*, consists of formal sums of monomials over x_1, \dots, x_n , where the set of coefficients is reduced by making all non-zero coefficients equivalent to 1. In other words, these are polynomials where the coefficients vary over the Boolean ring $\{0, 1\}$.

As an example, the sum of the *abstract* polynomial $x + y$ and the *abstract* polynomial $y + z$ is $x + y + z$, not $x + 2y + z$, as in the ring of coefficients, $1 + 1 = 1$ (also $1 \cdot 1 = 1$). The intention is that an abstract polynomial represents all the polynomials which are formed by varying its non-zero coefficients, thus $x + y$ represents $ax + by$ for all $a, b > 0$. We can also view an abstract polynomial as a set of monomials (whose coefficients are, implicitly, 1). We use AMP as an abbreviation for abstract multi-polynomial. We use $\alpha(\mathbf{p})$ for the abstraction of \mathbf{p} (obtained by modifying all non-zero coefficients to 1). Conversely, if \mathbf{p} is abstract we can “coerce” it to a concrete polynomial by considering the 1-valued coefficients as being literally 1; this is how one should read any statement that uses an abstract polynomial where a concrete function is needed. Abstract polynomials are used to ensure that the analysis is finite, since there is a finite number of different abstract polynomials of any given degree. They correspond to the use of big-O (or Θ) notation in typical algorithm analysis.

The set of abstract multi-polynomials is denoted by AMPo1 ; this has a composition operation $\mathbf{p} \bullet \mathbf{q}$, which relates to the standard composition $\mathbf{p} \circ \mathbf{q}$ by $\alpha(\mathbf{p}) \bullet \alpha(\mathbf{q}) = \alpha(\mathbf{p} \circ \mathbf{q})$; the different operator symbol (“ \bullet ” versus “ \circ ”) helps in disambiguating the meaning of an expression (referring to abstract polynomials versus concrete ones).

In the special case of *abstract univariate* polynomials, it is useful to restrict them to be monomials x^d ; this can be done since in “big Θ ” bounds, a univariate polynomial reduces to its leading monomial. We allow the monomial x^0 , as well as 0, and thus we obtain a semiring isomorphic to $\langle \mathbb{N} \cup \{-\infty\}, \max, + \rangle$.

Definition 2.4. A *univariate state* is an n -tuple of univariate monomials. A univariate state is called positive if it does not include a zero.

Based on this definition, a univariate state can be written as $\langle x^{d_1}, \dots, x^{d_n} \rangle$ and internally represented by either a vector of integers $\langle d_1, \dots, d_n \rangle$ (for a non-zero monomial) or the constant 0.

3. FORMAL STATEMENT OF THE PROBLEM AND OUR RESULTS

Our goal is to prove polynomial-space complexity bounds for the tight-bound problem. In [BAH20], our algorithm was encumbered by creating and calculating with exponentially-large objects. In particular, we argued that there are commands for which a set of multi-polynomials that provides tight bounds must include exponentially many elements. In order to reduce the size of objects *internal* to the algorithm, we have to rewrite the algorithm, but first we circumvent the issue of exponential *output size* by adjusting the specification of the algorithm. This can be done in two ways. First, as one commonly does in Complexity Theory, one can reduce the problem to a decision problem:

Problem 3.1. *Multivariate decision problem:* Given a core-language program P and a monomial $m(x_1, \dots, x_n) = x_1^{d_1} \dots x_n^{d_n}$, report whether m constitutes a “big-Omega” lower bound on the highest final value of \mathbf{X}_n for initial state $\langle x_1, \dots, x_n \rangle$.

There is, of course, no loss of generality in fixing the queried variable to be the highest-numbered; neither do we lose generality by restricting the bounds to monomials (a polynomial is an asymptotic lower bound if and only if all its monomials are). Note that we have chosen to focus on worst-case lower bounds, unlike most publications in this area, which focus on upper bounds for worst-case results. Our reader should keep in mind, though, that [BAH20] shows that *for our core language, whenever a variable is polynomially bounded, there is a set of polynomials which provides tight bounds for all executions*. This means that if we form the function “max of (all the lower bounds),” we have a tight worst-case upper bound. Thus solving for lower bounds also solves the upper bound problem. Furthermore, we note that if we are looking at a polynomially-bounded variable, the set of abstract monomials (or even polynomials) that *lower-bound the final value* is finite. This is convenient (for instance, the expression “max of (all the lower bounds)” can be explicitly written out, if desired). We refer to a polynomial in this set as *attainable* (since the program can compute values that reach, or surpass, this polynomial, asymptotically).

We call a function *attainable* (by a program) if its values can be reached or surpassed.

Definition 3.2 (Attainable). Consider a command \mathbf{C} , and a function $f : \mathbb{N}^n \rightarrow \mathbb{N}^n$ we say that f is *attainable* by \mathbf{C} if there are constants $d > 0$ and $\mathbf{x}_0 \in \mathbb{N}^n$ such that for all $\mathbf{x} \geq \mathbf{x}_0$ there is a \mathbf{y} such that

$$\mathbf{x}[\mathbf{C}]\mathbf{y} \text{ and } \mathbf{y} \geq df(\mathbf{x}).$$

Given a function $f' : \mathbb{N}^n \rightarrow \mathbb{N}$, extend the definition to f' by applying the above condition to $f(\mathbf{x}) = \langle 0, \dots, 0, f'(\mathbf{x}) \rangle$, in other words we use f' as a lower bound on a single variable (\mathbf{X}_n).

We remark that in the above inequation, d ranges over real numbers, and may be smaller than one.

Thus, for a program that starts with data x_1 and x_2 in its two variables, and non-deterministically chooses to exchange them, the bounds $\langle x_1, x_2 \rangle$ and $\langle x_2, x_1 \rangle$ are both attainable, while $\langle x_1, x_1 \rangle$ is not.

In the current work, we are attempting to find tight bounds only for variables that are polynomially bounded in terms of the input (the initial contents of the variables).

The problem of identifying which variables are polynomially bounded is completely solved in [BJK08], by a polynomial-time algorithm. We will tacitly rely on that algorithm. More precisely, we assume that the algorithm is invoked as a preprocessing step for the analysis of each loop in the program, and allows for excluding variables that may grow super-polynomially in the given loop. This reduces the problem of handling *any* loop in our language to that of handling loops in which all variables are polynomially bounded; the key observation here is that, due to the restricted arithmetics in our language, which only include addition and product, any value derived from a super-polynomial value is also super-polynomial. Hence, *ignoring* variables which are not polynomially bounded does not lose any information necessary for analyzing the polynomially bounded ones.

Note that in the formulation as a decision problem, the user queries about a given polynomial; we think that it is a natural use-case, but one can also ask the analyzer to furnish a bound:

Problem 3.3. *Multivariate bound generation problem:* Given a core-language program P , and an integer d , find a polynomial $p(x_1, \dots, x_n)$ with the degree in every variable bounded by d , such that p constitutes a “big-Omega” lower bound on the highest final value of X_n as a function of x_1, \dots, x_n .

The function of the input parameter d is to (possibly) reduce the complexity of the search—not looking further than what the user is wishing to find. There is, of course, no loss of generality in fixing the queried variable to be the highest-numbered.

For this problem we claim a solution whose space complexity is polynomial in the input size, namely the program size plus the bit-size of d . Recall that if we have a polynomial-space non-deterministic algorithm, we can also determinize it in polynomial space (this is the essence of the famous equality $\text{NPSpace} = \text{PSPACE}$, which follows from Savitch’s Theorem [Jon97]). This means that it is possible to transform the non-deterministic solution to Problem 3.3 (as long as it is complete) into a deterministic generator that writes out one monomial at a time. Moreover, by exhaustive search it is possible to generate all the monomials needed to bound all program executions, even when their number is exponential.

It should thus be clear that in a similar manner we can also search for a bound of highest degree d_{max} . The complexity of our solution to the multivariate-bound problem (Sections 11 and 12) is *output-sensitive* in the sense that the bound is polynomial in the program size and the bit-size of d_{max} .

A main idea in this paper is to reduce Problem 3.1 to a problem concerning univariate bounds. In this problem, we assume that we are given a *univariate initial state* in terms of a single input variable x . The initial state assigns to every variable X_i an initial symbolic value of form x^{d_i} for some integer $d_i \geq 0$, e.g., $\mathbf{i}(x) = \langle x, x^3, x^2 \rangle$. The case $d_i = 0$ intuitively represents a variable whose initial value should be treated as an unknown constant (not dependent on x , but possibly large). The univariate form seems to be rather common in textbook examples of algorithm analysis (where the single parameter in terms of which we express the bounds is usually called n), and the facility of stating initial polynomial values for multiple variables may be useful to express questions like: *An algorithm processes a graph. Its running time depends on the number of vertices n and the number of edges m . Bound its worst-case time, in terms of n , given that $m \leq n^2$.* In this case we would use an initial state like $\langle n, n^2 \rangle$. Our goal is now to find an attainable univariate state, providing a lower bound on the worst-case results of a computation from the given initial state. Here is the formal definition; note that we are providing simultaneous bounds for all variables.

Definition 3.4. Given a command \mathcal{C} , a function $g : \mathbb{N} \rightarrow \mathbb{N}^n$ and an initial univariate state \mathbf{i} , we say that g is attainable (by \mathcal{C} , given \mathbf{i}) if there are constants $d > 0$, x_0 such that for all $x \geq x_0$ there is a \mathbf{y} such that

$$\mathbf{i}(x) \llbracket \mathcal{C} \rrbracket \mathbf{y} \text{ and } \mathbf{y} \geq dg(x).$$

As in Definition 3.2, a function with codomain \mathbb{N} may be used and interpreted as a lower bound for just X_n .

We thus focus on the following problem.

Problem 3.5. *Univariate decision problem:* Given core-language program P , a positive² initial state \mathbf{i} , the index of a chosen variable X_j and an integer d , decide whether there is an attainable univariate state whose degree in the j th component is at least d .

We also solve the corresponding bound generation problem—even for simultaneous bounds, namely

Problem 3.6. *Univariate bound generation problem:* Given a core-language program P , a positive initial state \mathbf{i} , a variable index j and an integer d , non-deterministically find an attainable univariate state whose degree in the j th component is at least d .

We have included the parameter d in this problem formulation as well, since we will be able to bound the complexity of our algorithm in terms of d (instead of d_{max}).

Here is a summary of the results we are going to prove regarding the above problems:

- In Section 7, we give a polynomial space algorithm to solve Problem 3.6.
- The algorithm in Section 7, along with the complexity analysis in Section 8, prove that Problem 3.5 and Problem 3.6 can be solved in polynomial space. Note that this means that the space complexity is polynomial in the input size, consisting of the size of program P and the bit-size of the numbers in \mathbf{i} and d .
- In Section 10 we show that this result is essentially optimal since the problem is PSPACE-hard. This clearly means that the more involved problems (for multivariate bounds) are also PSPACE-hard.
- In Section 11 we show a reduction of the multivariate-bound decision problem (Problem 3.1) to the univariate problem (Problem 3.5), and obtain in this way a solution with space complexity polynomial in terms of the size of the program and the maximum degree.
- In Section 12 we solve the multivariate-bound generation problem (Problem 3.3).

We'd like to mention that solving the above problems also solves some variants that can be reduced to tight-bound computation:

- We can find tight bounds on the number of *visits* to a given set of program locations (counting all locations gives the program's time complexity, in the ordinary unit-cost model). The reduction is to instrument the program with a counter that counts the visits to the locations of interest [BAH20, Section 2].
- Sometimes we want to verify a bound for the values of a variable *throughout* the program, not just at its end. This can also be easily reduced to our problem by instrumentation.
- Similarly, one can find a tight bound on the highest value assumed by a variable at a given program location.

²A zero in the initial state can be accommodated in an extension discussed in Section 13, where an assignment $X := 0$ is also allowed.

4. PROPERTIES OF THE CORE LANGUAGE

We next present some important properties of the core language that follow from [BAH20] and are required for our result. First, let us cite the main result of [BAH20]:

Theorem 4.1 ([BAH20]). *There is an algorithm which, for a command \mathbf{C} , over variables X_1 through X_n , outputs a set \mathcal{B} of multi-polynomials, such that the following hold, where PB is the set of indices i of variables X_i which are polynomially bounded under $\llbracket \mathbf{C} \rrbracket$.*

(1) (*Bounding*) *There is a constant c such that*

$$\forall \mathbf{x}, \mathbf{y} . \mathbf{x} \llbracket \mathbf{C} \rrbracket \mathbf{y} \implies \exists \mathbf{p} \in \mathcal{B} . \forall i \in PB . y_i \leq c \mathbf{p}[i](\mathbf{x})$$

(2) (*Tightness*) *There are constants $d > 0$, \mathbf{x}_0 such that for every $\mathbf{p} \in \mathcal{B}$, for all $\mathbf{x} \geq \mathbf{x}_0$ there is a \mathbf{y} such that*

$$\mathbf{x} \llbracket \mathbf{C} \rrbracket \mathbf{y} \text{ and } \forall i \in PB . y_i \geq d \mathbf{p}[i](\mathbf{x}).$$

We rely on this result because our new algorithm will be proven correct by showing that it matches the bounds of the previous algorithm. But the implied *existence result* is crucial in itself: the fact that a set \mathcal{B} as above exists, which provides both upper bounds (clause “Bounding”) and matching worst-case lower bounds (clause “Tightness”). This property is tied to the limitations of our language: in a Turing-complete programming language this clearly does not hold—some times no polynomial tight bound exists, as for a program that inputs x and computes $\log x$, \sqrt{x} , etc. In other cases the bound is only “tight” in the sense that it is matched by a lower bound for infinitely many inputs, but not for all inputs (this is sometimes expressed by the notation Ω_∞); e.g., a program that squares a number only if it is odd. These cases cannot happen with our language.

Since a polynomial is a sum of a constant number of monomials, by simple arithmetics we obtain:

Corollary 4.2. *For a command \mathbf{C} , over variables X_1 through X_n , assume that X_n is polynomially bounded in terms of the variables’ initial values. Then there is a set \mathcal{S} of monomials such that $\max \mathcal{S}$ is a tight asymptotic bound on the highest value obtained by X_n at the conclusion of \mathbf{C} ; or in more detail,*

(1) (*Bounding*) *There is a constant c such that*

$$\forall \mathbf{x}, \mathbf{y} . \mathbf{x} \llbracket \mathbf{C} \rrbracket \mathbf{y} \implies \exists m \in \mathcal{S} . y_n \leq c m(\mathbf{x})$$

(2) (*Tightness*) *There are constants d , \mathbf{x}_0 such that for every $m \in \mathcal{S}$, for all $\mathbf{x} \geq \mathbf{x}_0$ there is a \mathbf{y} with*

$$\mathbf{x} \llbracket \mathbf{C} \rrbracket \mathbf{y} \text{ and } y_n \geq d m(\mathbf{x}).$$

We refer to monomials satisfying the Tightness condition as attainable; they constitute valid answers to Problem 3.3 (subject to the degree bound); and the fact that they also provide an upper bound is, again, an important property that depends on the limitations of the programming language. This will be used in Sections 11 and 12 where we check or compute multivariate bounds in the form of monomials.

5. ALGORITHMIC IDEAS

In this section we give a brief overview and motivation of key elements in our solution. In particular, we relate these elements to prior work. We discuss symbolic evaluation, the motivation for working with univariate bounds, and data-flow matrices. This is followed by an overview of the way we prove correctness of this algorithm, and the way we finally solve the multivariate-bound problem.

5.1. Symbolic evaluation. This is an old concept. We interpret the program, with states mapping variables not to integers, but to polynomials in x_1, \dots, x_n , where x_i represents the initial value of the corresponding variable X_i . Due to the restricted arithmetics in our language, we can carry out symbolic evaluation precisely, staying within the domain of multivariate polynomials. Thus a program state is a MP. However, when there are different program paths that reach the same point, we can get different polynomials, and therefore we construct *sets* of polynomials. The abstraction to *abstract polynomials* may reduce the size of this set. Here is an example of a piece of program code, along with a symbolic state, as a set of AMPs, before and after each command:

$$\begin{aligned} & \{ \langle x_1, x_2, x_3, x_4 \rangle \} \\ X_2 & := X_2 + X_4; \\ & \{ \langle x_1, x_2 + x_4, x_3, x_4 \rangle \} \\ X_4 & := X_3; \\ & \{ \langle x_1, x_2 + x_4, x_3, x_3 \rangle \} \\ \text{choose } & \{ X_1 := X_1 + X_3 \} \text{ or } \{ X_2 := X_2 + X_3 \} \\ & \{ \langle x_1 + x_3, x_2 + x_4, x_3, x_3 \rangle, \langle x_1, x_2 + x_3 + x_4, x_3, x_3 \rangle \} \end{aligned}$$

This is all quite straight-forward until one comes to a loop. Since we are analysing the program statically, we have to take into account any number of iterations. Consequently, had we worked with concrete polynomials, they would grow indefinitely.

Example 5.1. Consider this loop:

```
loop X2 {
  X1 := X1 + X2;
}
```

Evaluating with *concrete* polynomials, the value for X_1 would take the values $x_1, x_1 + x_2, x_1 + 2x_2, x_1 + 3x_2$, etc. Working with abstract polynomials, the coefficients are ignored, and we have a finite set, which is reached after a finite number of iterations. In our example, the value for X_1 does not rise beyond $x_1 + x_2$. This is good, but unfortunately, we also lose completeness: $x_1 + x_2$ is not a valid asymptotic upper bound on the final value of X_1 ; a correct upper bound is $x_1 + x_2^2$, taking into account the bound x_2 on the number of loop iterations. The passage from $x_1 + x_2$ to $x_1 + x_2^2$ is based on correctly *generalizing* the behaviour that we see in one iteration to the long-term behaviour, and deducing that the *increment* $+x_2$ may be applied to this variable at most x_2 times. Making such generalizations precisely is the main challenge addressed in the algorithm of [BAH20].

Our algorithm in the cited work gave a precise set of bounds but suffered from efficiency problems. This is inherent in the approach of generating symbolic polynomials. In fact, it

is easy to write pieces of code that generate exponentially big sets of expressions—try the following:

```

choose { X3 := X1 } or { X3 := X2 } ;
choose { X4 := X1 } or { X4 := X2 } ;
. . .
choose { Xn := X1 } or { Xn := X2 }

```

5.2. Univariate bounds. Suppose that we decide to compute only univariate bounds. We have a single input parameter x and we want to express everything in terms of x . Then, if we also ignore coefficients, the only functions we need are x, x^2, x^3, \dots and the efficiency issue, at least the problem of big expressions, is resolved. An abstract state is just an n -tuple of degrees and takes polynomial space. However, important information is lost. If we start with all variables set to x , and the loop body sets $X_1 := X_1 + X_2$, we get the symbolic value $2x$ in X_1 after the first iteration (or just x if we ignore coefficients). How do we know that this becomes $x + tx$ after t iterations? We cannot deduce it from the final expression. We are compelled to also keep track of how the value was computed—so we know that there is an increment that can accumulate upon iteration. Note that the fact that the value after one iteration is different from the value before does not necessarily mean that the difference will accumulate—consider a loop body that simply sets $X_1 := X_2$, where the initial value of X_2 is bigger than that of X_1 .

5.3. Data-flow matrices. The discussion in the last paragraph hints that it may be useful to record data-flow: does x'_1 (value of X_1 after an iteration) depend on x_1 (its initial prior to the iteration)? On x_2 ? On both? Such information may be represented by a *bipartite graph* with arcs from x_1, x_2, \dots, x_n leading to x'_1, x'_2, \dots, x'_n and showing how values propagate, or a *data-flow matrix*, which is the same information in matrix form. This is a very concise representation which does not record the actual expressions computed, only the Boolean property that some x'_j depends, or does not depend, on x_i . Some of the previous results in complexity analysis [KA80, KN04] showed that the existence of polynomial bounds on computed values may sometimes be deduced by examining the dependence graph. Later works [NW06, JK09] showed that by enriching the information in data-flow matrices (allowing for a finite number of “dependence types”) one has sufficient information to soundly conclude that a result is polynomially bounded in a larger set of programs. As a basic example, the technique allows us to distinguish a loop body that sets $x'_1 = 2x_1$ (doubling the value of X_1 , leading to exponential growth upon iteration) from one that sets $x'_1 = x_1 + x_2$ (which entails polynomial growth). Ben-Amram et al. [BJK08] derived a complete decision procedure for polynomial growth-rates in the Core Language, still by tracking data-flow properties (with an abstraction similar to data-flow matrices) and without explicitly computing any bounds. In the current work, we use data-flow matrices in addition to a symbolic state, to aid the analysis of loops, specifically the precise generalization of iterations that seem to increase certain variables. Surprisingly, it turns out that all we need is Boolean matrices which record only linear data-dependence.

5.4. Correctness. We have motivated our work by wishing to solve the efficiency issue with the algorithm of [BAH20]. However that algorithm—the so-called *Closure Algorithm*—remains the foundation of our current contribution. We both motivate the polynomial-space algorithm, and formally prove its correctness, based on the idea that it should match the results of the Closure Algorithm, in the following sense: the Closure Algorithm computes AMPs; for instance, for the loop in Example 5.1, it outputs the AMP $\mathbf{p} = \langle x_1 + x_2^2, x_2 \rangle$. Our algorithm works with an initial univariate state, e.g., $\mathbf{x} = \langle x^3, x^2 \rangle$. An abstract multi-polynomial can be applied to a univariate state, e.g., $\mathbf{p} \bullet \mathbf{x} = \langle x^4, x^2 \rangle$ (as already mentioned, we reduce abstract univariate polynomials to their leading monomial). The correctness assertion which we will prove for the polynomial-space algorithm is that it can match the result $\mathbf{p} \bullet \mathbf{x}$ (though it never maintains \mathbf{p}). Our algorithm will be non-deterministic, and in this example it will also generate additional final abstract states, including $\langle x^3, x^2 \rangle$. In the *soundness* part of the proof, we argue that every such result is attainable, by showing that it is bounded from above by $\mathbf{p} \bullet \mathbf{x}$ for \mathbf{p} computed by the Closure Algorithm; in the *completeness* part, we show that for every \mathbf{p} computed by the Closure Algorithm, we can reach (or surpass) $\mathbf{p} \bullet \mathbf{x}$ (clearly, if $\mathbf{p} \bullet \mathbf{x}$ is a maximal element in the set of attainable bounds, then based on soundness, we will reach but not exceed it). Thus, our proof does not validate the new analysis by relating its results with the concrete semantics of the analysed program, but with the results of the Closure Algorithm, which [BAH20] already related to concrete results.

5.5. Multivariate bounds. As the restriction to univariate bounds was a key element in the development and proof of the polynomial-space algorithm, it came as a surprise that the same algorithm (practically the same code)—with a change of domain from univariate AMPs to multivariate ones—can compute tight multivariate bounds. But our correctness proof, as outlined above, *does not* work for the multivariate algorithm (otherwise we could have skipped the univariate stage altogether). We arrive at the result for multivariate bounds by first establishing a *reduction* among decision problems—the attainability of a multivariate bound and that of univariate bounds. We study the degree vectors (d_1, d_2, \dots, d_n) of monomial bounds $x_1^{d_1} \dots x_n^{d_n}$, and show that these vectors define a polyhedron in n -space, and that a linear-programming view of the problem allows us to decide membership in this polyhedron using the univariate analysis. This linear-programming approach also shows that in order to obtain a tight upper bound it is not necessary to take a **max** of all the attainable bounds—it suffices to use those that figure as the vertices of the polyhedron. We proceed to use this insight, along with the linear-programming connection to univariate bounds, to prove that we can compute a sound and complete set of multivariate bounds—where by *sound* we mean that the bounds are attainable, i.e., constitute lower bounds on some possible results, and by *complete* we mean that their maximum is an upper bound on *all* possible results.

6. THE CLOSURE ALGORITHM

In this section, we present a version of the algorithm of [BAH20], computing in exponential space and time. We call it “the Closure Algorithm” since an important component of the algorithm (and a cause of combinatorial explosion) is the computation of transitive closure to find the effect of any number of loop iterations. We need the Closure Algorithm because

our polynomial-space solution evolved from this one, and moreover our *proof* for the efficient algorithm relies on the correctness of its predecessor.

In fact, the algorithm below is already a step beyond [BAH20]: even if it is still exponential, it is somewhat simplified. Since this is not the main contribution here, we have decided to present it, in this section, concisely and without proofs, to allow the reader to understand and proceed quickly to our new polynomial-space algorithm. For completeness, Appendix A includes proofs that show that the Closure Algorithm below is equivalent to the old one [BAH20].

6.1. Symbolic semantics. We present our analysis as symbolic semantics that assigns to every command C a symbolic abstract value $\llbracket C \rrbracket^S \in \wp(\mathbf{AMPo1})$, that is, a set of AMPs. For atomic commands, this set is a singleton:

$$\begin{aligned} \llbracket \mathbf{skip} \rrbracket^S &= \{Id\} & \llbracket X_i := X_j \rrbracket^S &= \{\mathbf{set}_{ij}\} \\ \llbracket X_i := X_j + X_k \rrbracket^S &= \{\mathbf{add}_{ijk}\} & \llbracket X_i := X_j * X_k \rrbracket^S &= \{\mathbf{mul}_{ijk}\} \end{aligned}$$

where \mathbf{set}_{ij} , \mathbf{add}_{ijk} and \mathbf{mul}_{ijk} are AMPs defined by:

$$\begin{aligned} \mathbf{set}_{ij}(x_1, \dots, x_n) &= \langle x_1, \dots, x_{i-1}, x_j, x_{i+1}, \dots, x_n \rangle \\ \mathbf{add}_{ijk}(x_1, \dots, x_n) &= \langle x_1, \dots, x_{i-1}, x_j + x_k, x_{i+1}, \dots, x_n \rangle \\ \mathbf{mul}_{ijk}(x_1, \dots, x_n) &= \langle x_1, \dots, x_{i-1}, x_j x_k, x_{i+1}, \dots, x_n \rangle. \end{aligned}$$

For composite commands, except the loop, the definition is also straightforward:

$$\begin{aligned} \llbracket \mathbf{choose} C_1 \text{ or } C_2 \rrbracket^S &= \llbracket C_1 \rrbracket^S \cup \llbracket C_2 \rrbracket^S \\ \llbracket C_1; C_2 \rrbracket^S &= \llbracket C_2 \rrbracket^S \bullet \llbracket C_1 \rrbracket^S. \end{aligned}$$

To handle a loop command, $\mathbf{loop} X_\ell \{C\}$, we first compute $\mathcal{S} = \llbracket C \rrbracket^S$, obtaining a representation of the possible effects of any *single iteration*. Then we have to apply certain operations to calculate, from \mathcal{S} , the effect of the entire loop:

$$\llbracket \mathbf{loop} X_\ell \{C\} \rrbracket^S = LC(\llbracket C \rrbracket^S)[x_\ell/\tau].$$

The rest of this section builds up to the definition of the function LC and the explanation of the above expression.

Example 6.1. Consider the following loop (also considered in Section 5):

```
loop X3 {
  X2 := X2 + X4;
  X4 := X3;
  choose { X1 := X1 + X3 } or { X2 := X2 + X3 }
}
```

The abstraction of the *loop body* is the following set of AMPs:

$$\mathcal{S} = \{\langle x_1 + x_3, x_2 + x_4, x_3, x_3 \rangle, \langle x_1, x_2 + x_3 + x_4, x_3, x_3 \rangle\}.$$

We will later show how this is used to compute the effect of the whole loop.

6.2. Simple Disjunctive Loops, Closure and Generalization. We cite some definitions from [BAH20] that we use in presenting the inference of the effect of a loop from an abstraction of its body.

Definition 6.2. A *polynomial transition (PT)* represents the passage from a state $\mathbf{x} = \langle x_1, \dots, x_n \rangle$ to a new state $\mathbf{x}' = \langle x'_1, \dots, x'_n \rangle = \mathbf{p}(\mathbf{x})$ where $\mathbf{p} = \langle \mathbf{p}[1], \dots, \mathbf{p}[n] \rangle$ is a multi-polynomial.

Definition 6.3. A *simple disjunctive loop (SDL)* is a finite set \mathcal{S} of AMPs, representing polynomial transitions.

The loop is “disjunctive” because the meaning is that, in every iteration, any of the given transitions may be chosen. The sense in which *abstract* MPs represent *concrete* transitions calls for a detailed definition but we omit it here, since in our context the intent should be clear: we expect \mathcal{S} to be the result of analyzing the loop body with an analysis that generates asymptotically tight bounds. Importantly, a SDL does not specify the number of iterations; our analysis of a SDL generates results parameterized by the number of iterations as well as the initial state. For this purpose, we now introduce τ -polynomials, where τ (for “time”) is a parameter to represent the number of iterations.

Definition 6.4. τ -polynomials are polynomials in x_1, \dots, x_n and τ .

As τ represents the number of iterations, it is not a component of the state vector, which still has dimension n . If p is a τ -polynomial, then $p(v_1, \dots, v_n)$ is the result of substituting each v_i for the respective x_i ; and we also write $p(v_1, \dots, v_n, t)$ for the result of substituting t for τ as well.

We form multi-polynomials from τ -polynomials to represent the effect of a variable number of iterations. For example, the τ -polynomial transition $\langle x'_1, x'_2 \rangle = \langle x_1, x_2 + \tau x_1 \rangle$ represents the effect of repeating (τ times) the assignment $\mathbf{X}_2 := \mathbf{X}_2 + \mathbf{X}_1$. The effect of iterating the composite command: $\mathbf{X}_2 := \mathbf{X}_2 + \mathbf{X}_1$; $\mathbf{X}_3 := \mathbf{X}_3 + \mathbf{X}_2$ has an effect described by $\mathbf{x}' = \langle x_1, x_2 + \tau x_1, x_3 + \tau x_2 + \tau^2 x_1 \rangle$ (note that this is an upper bound which is not reached precisely, but is correct up to a constant factor). We denote the set of τ -multi-polynomials by $\tau\mathbf{MPol}$. We should note that composition $\mathbf{q} \circ \mathbf{p}$ over $\tau\mathbf{MPol}$ is performed by substituting $\mathbf{p}[i]$ for each occurrence of x_i in \mathbf{q} . Occurrences of τ are unaffected (since τ is not part of the state).

The notion of *abstract (multi-) polynomials* is extended naturally to abstract τ -(multi-) polynomials. We denote the set of abstract τ -polynomials (respectively, multi-polynomials) by $\tau\mathbf{APol}$ (respectively, $\tau\mathbf{AMPol}$). The *SDL Problem* is to compute, from \mathcal{S} , a set of abstract τ -multi-polynomials that represent the effect of any number of loop transitions, where τ is used to express dependence on the number of transitions. As for general programs, we focus on loops which are polynomially bounded, and search for attainable bounds; these notions are defined as follows.

Definition 6.5. A SDL \mathcal{S} is said to be *polynomially bounded* when there exists a τ -MP, \mathbf{b} , such that for all $\mathbf{x} \in \mathbb{N}^n$, if we start with \mathbf{x} and consecutively apply t arbitrary transitions from \mathcal{S} , the final state \mathbf{y} satisfies $\mathbf{y} \leq \mathbf{b}(\mathbf{x}, t)$.

Note that we require a polynomial bound dependent on the number of iterations t , considering t as an independent variable. This makes the SDL analysis independent of the program context where the SDL is extracted from. Clearly, when we plug the results of SDL analysis back into a context where the number of iterations is also bounded by a polynomial,

we shall obtain a polynomial bound on the final state in terms of the program variables (eliminating t).

Definition 6.6. Given SDL \mathcal{S} and a function $\mathbf{f} : \mathbb{N}^{n+1} \rightarrow \mathbb{N}^n$ (in the current context, \mathbf{f} will be a τ -MP) we say that \mathbf{f} is *attainable over \mathcal{S}* if there are constants $d > 0$, \mathbf{x}_0 such that for all $\mathbf{x} \geq \mathbf{x}_0$, for infinitely many values $t > 0$, there exist $\mathbf{y}_0, \dots, \mathbf{y}_t$ such that

$$\mathbf{y}_0 = \mathbf{x}; \quad \forall i < t. \exists \mathbf{p} \in \mathcal{S}. \mathbf{y}_{i+1} = \mathbf{p}(\mathbf{y}_i); \quad \text{and } \mathbf{y}_t \geq d\mathbf{f}(\mathbf{x}, t).$$

We remark that SDLs enjoy the properties we pointed out in Section 4. In particular, a complete set of attainable lower bounds provides a tight asymptotic upper bound as well. In [BAH20], we studied the SDL problem, and what we present next is an improved version of our solution from that article; the difference is explained at the end of this section. The solution consists of applying two operators to a set of (τ)-AMPs, which we shall now define. The first, *abstract closure*, naturally represents the limit-set of accumulated iterations.

Definition 6.7 (abstract closure). For finite $P \subset \tau\text{AMPo1}$, we define $Cl(P)$ to be the smallest set including Id and the elements of P , and closed under AMP composition.

In [BAH20] we proved that when loop \mathcal{S} is polynomially bounded, the abstract closure of \mathcal{S} is finite, and can be computed in a finite number of steps. Briefly, the reason that the closure must be finite, when the loop is polynomially bounded, is that the polynomials in the abstract closure are of bounded degree (since they are valid lower bounds). Moreover, there is a finite number of different abstract polynomials of any given degree.

Example 6.8. In Example 6.1 we have seen that

$$\mathcal{S} = \{ \langle x_1 + x_3, x_2 + x_4, x_3, x_3 \rangle, \quad \langle x_1, x_2 + x_3 + x_4, x_3, x_3 \rangle \}.$$

Computing $Cl(\mathcal{S})$ adds Id as well as the composition of the above two AMPs, which equals

$$\mathbf{p} = \langle x_1 + x_3, x_2 + x_3 + x_4, x_3, x_3 \rangle.$$

The reader is invited to verify that this gives a composition-closed set.

The second operation is called *generalization* and its role is to capture the behaviour of variables that grow by accumulating increments in the loop, and make explicit the dependence on the number of iterations. The identification of which additive terms in a MP should be considered as increments that accumulate is at the heart of our problem, and its solution led to the definition of *iterative MPs* and *iterative kernel* below.

Definition 6.9. The *support* $\text{sup } p$ of a polynomial p is the set of variables on which it depends, identified by their indices, e.g., $\text{sup}(x_1x_3 + x_4) = \{1, 3, 4\}$.

Definition 6.10. For \mathbf{p} an (abstract) multi-polynomial, we say that x_i is *self-dependent* in \mathbf{p} if $i \in \text{sup } \mathbf{p}[i]$. We also say that the entry $\mathbf{p}[i]$ is self-dependent; the choice of term depends on context and the meaning should be clear either way. We denote by $\text{SD}(\mathbf{p})$ the set $\{i : i \in \text{sup } \mathbf{p}[i]\}$, i.e., the self-dependent variables of \mathbf{p} .

It is easy to see that, given that variable X_i is polynomially bounded in a loop, if x_i is self-dependent, then $\mathbf{p}[i]$ must be of the form $x_i + q$, where q is a polynomial that does not depend on x_i (we say that x_i is *linearly self-dependent* in \mathbf{p}). Otherwise, \mathbf{p} represents a transition that multiplies X_i by a non-constant quantity and iterating \mathbf{p} will cause exponential growth.

Definition 6.11. We call an (abstract) MP \mathbf{p} *doubling-free* if for any i , if entry $\mathbf{p}[i]$ has a monomial $\mathbf{m} \cdot x_i$ then $\mathbf{m} = 1$.

Assuming that we have made sure that a loop under analysis is polynomially bounded (as briefly discussed in Section 3 and more fully in [BAH20]), then all the loop transitions must be doubling-free. Hence if $\mathbf{p}[i]$ depends on x_i , it must include the monomial x_i .

Definition 6.12. Let $\mathbf{p} \in \text{AMPo1}$. A monomial $x_{i_1} \dots x_{i_n}$ is called *iterative* with respect to \mathbf{p} if all i_k belong to $\text{SD}(\mathbf{p})$. We call \mathbf{p} *iterative* if all its monomials are iterative with respect to \mathbf{p} .

Iterative AMPs are crucial in the analysis of loops as, intuitively, they represent a transformation that can happen multiple times. To see what we mean compare the following assignment statements:

- $\mathbf{X}_1 := \mathbf{X}_2 + \mathbf{X}_3$, represented by the AMP $\mathbf{p} = \langle x_2 + x_3, x_2, x_3 \rangle$. Both x_2 and x_3 are self-dependent and therefore \mathbf{p} is iterative. In fact, if iterated any number of times, the resulting AMP is the same (in this case even the concrete MP remains the same, since there is no growth in the loop).
- $\mathbf{X}_1 := \mathbf{X}_1 + \mathbf{X}_3$, represented by the AMP $\mathbf{p} = \langle x_1 + x_3, x_2, x_3 \rangle$. All three variables are self-dependent and therefore \mathbf{p} is iterative. In fact, if iterated any number of times, the resulting AMP is the same. Importantly, the concrete MP will not be the same: increments of x_3 *accumulate* in variable \mathbf{X}_1 . The algorithm will have to correctly express this growth, generalizing \mathbf{p} to $\langle x_1 + \tau x_3, x_2, x_3 \rangle$.
- $\mathbf{X}_1 := \mathbf{X}_1 + \mathbf{X}_3; \mathbf{X}_3 := \mathbf{X}_2$, represented by the AMP $\mathbf{q} = \langle x_1 + x_3, x_2, x_2 \rangle$. Here x_1 and x_2 are self-dependent, but x_3 (used in $\mathbf{q}[1]$) is not, therefore \mathbf{q} is not iterative. In fact, if iterated twice, we get $x'_1 = x_1 + x_3 + x_2$; informally, the action of the first application is to add x_3 to \mathbf{X}_1 while the action of the second is to add x_2 . It would have been incorrect to generalize from the first step and assume that increments of x_3 will be accumulated on iteration.

Definition 6.13 (ordering of multi-polynomials). For $\mathbf{p}, \mathbf{q} \in \text{AMPo1}$ we define $\mathbf{p} \sqsubseteq \mathbf{q}$ to hold if every monomial of $\mathbf{p}[i]$ also appears in $\mathbf{q}[i]$ for $i = 1, \dots, n$. We then say that \mathbf{p} is a *fragment* of \mathbf{q} .

Lemma 6.14. Let $\mathbf{p} \in \text{AMPo1}$. Then \mathbf{p} has a unique maximal iterative fragment, which shall be denoted by (\mathbf{p}) (“the iterative kernel of \mathbf{p} ”).

Proof. Every monomial which is not iterative in \mathbf{p} will not be iterative in any fragment of \mathbf{p} , so form \mathbf{q} by deleting all these monomials. Then we must have $\mathbf{i} \sqsubseteq \mathbf{q}$ for any iterative $\mathbf{i} \sqsubseteq \mathbf{p}$. On the other hand, \mathbf{q} is clearly iterative, so \mathbf{q} is the maximal iterative fragment. \square

Note that even if all entries of \mathbf{p} are non-zero, (\mathbf{p}) may include zeros, e.g., for $\mathbf{p} = \langle x_3, x_2, x_1 + x_2 \rangle$, where $(\mathbf{p}) = \langle 0, x_2, x_2 \rangle$.

Consider a loop that has a transition \mathbf{p} with $\mathbf{p}[3] = x_3 + x_1x_2$, where x_1 and x_2 are also self-dependent in \mathbf{p} . Then the initial contents of variables \mathbf{X}_1 and \mathbf{X}_2 are preserved (or may even grow) when we iterate \mathbf{p} ; and \mathbf{X}_3 accumulates an increment of (at least) x_1x_2 on each iteration. This motivates the following definition.

Definition 6.15 (generalization). Let $\mathbf{p} \in \text{AMPo1}$ be iterative; we define the *generalization* of \mathbf{p} to be

$$\mathbf{p}^\tau[i] = \begin{cases} x_i + \tau q & \text{if } \mathbf{p}[i] = x_i + q \\ \mathbf{p}[i] & \text{otherwise.} \end{cases}$$

To continue the example presented above the definition, let

$$\mathbf{p} = \langle x_1, x_2 + x_1, x_3 + x_1x_2 \rangle$$

then

$$\mathbf{p}^\tau = \langle x_1, x_2 + \tau x_1, x_3 + \tau x_1x_2 \rangle.$$

We now define the operators used in analysing a loop command:

$$Gen(\mathcal{S}) = \mathcal{S} \cup \{(\mathbf{p})^\tau \mid \mathbf{p} \in \mathcal{S}\}, \quad (6.1)$$

$$LC(\mathcal{S}) = Cl(Gen(Cl(\mathcal{S}))). \quad (6.2)$$

Recall that we have defined

$$\llbracket \text{loop } \mathbf{x}_\ell \ \{\mathbf{C}\} \rrbracket^{\mathcal{S}} = LC(\llbracket \mathbf{C} \rrbracket^{\mathcal{S}})[x_\ell/\tau].$$

This means that one takes the set \mathcal{S} representing the analysis of the loop body, applies the function LC , which generates τ -AMPs, and concludes by substituting x_ℓ , the maximum number of iterations, for τ .

Example 6.16. In Example 6.1 we have seen that

$$\mathcal{S} = \{ \langle x_1 + x_3, x_2 + x_4, x_3, x_3 \rangle, \langle x_1, x_2 + x_3 + x_4, x_3, x_3 \rangle \}.$$

and in Example 6.8 we have seen that $Cl(\mathcal{S}) = \mathcal{S} \cup \{Id, \mathbf{p}\}$ where

$$\mathbf{p} = \langle x_1 + x_3, x_2 + x_3 + x_4, x_3, x_3 \rangle.$$

Next, we apply generalization. We will take a shortcut and apply it only to \mathbf{p} , as it subsumes the two previous AMPs. We have:

$$(\mathbf{p}) = \langle x_1 + x_3, x_2 + x_3, x_3, x_3 \rangle, \quad (6.3)$$

$$(\mathbf{p})^\tau = \langle x_1 + \tau x_3, x_2 + \tau x_3, x_3, x_3 \rangle, \quad (6.4)$$

Then, in the final closure computation, we construct the composition

$$(\mathbf{p})^\tau \bullet \mathbf{p} = \langle x_1 + x_3 + \tau x_3, x_2 + x_3 + x_4 + \tau x_3, x_3, x_3 \rangle$$

and this is as high as we get. So we finish by substituting x_3 (being the loop bound) for τ , obtaining the result

$$\langle x_1 + x_3 + x_3^2, x_2 + x_3 + x_4 + x_3^2, x_3, x_3 \rangle.$$

We conclude this section with some comments about the improvement of the above algorithm on [BAH20]. First, in [BAH20], iterative kernels were not used. Generalization was applied to *idempotent* elements in the closure set (satisfying $\mathbf{p} \bullet \mathbf{p} = \mathbf{p}$); an idempotent AMP is not necessarily iterative, and the definition of the generalization operator was slightly more complex than the definition for iterative AMPs. The choice to focus on idempotent elements was guided by an algebraic theorem that we used to prove the algorithm correct (Simon's Forest Factorization Theorem). Another difference is that in the original algorithm, the computation of the closure and subsequent generalization had to be performed in several rounds (until the set stabilized); in contrast, when generalization is applied to iterative kernels as done here, it suffices to compute closure, generalize and close once more, as expressed by (6.2).

7. THE POLYNOMIAL-SPACE ALGORITHM

In this section we provide the polynomial-space algorithm for the univariate bound problem. We will try to present the algorithm in a way which motivates it and clarifies why it works. First, let us remind the reader that we are trying to compute attainable functions: such a function describes a bound that can be reached (for almost all inputs, on at least one possible computation for every such input), and therefore provides a lower bound on the worst-case result; moreover, we know by Theorem 4.1 that a complete set of attainable bounds exists, in the sense that it also implies a tight upper bound (specifically, the **max** of this finite set of functions is a worst-case upper bound). Intuitively, since an attainable function is “something that can happen,” to find such functions we just have to evaluate the program symbolically and see what it computes: expressed as a function of the input, this will be an attainable function. Our evaluator will be non-deterministic and so generate a set of attainable functions. It will be then necessary to prove that this is a complete set.

7.1. The Simple Case: Without Additions in Loops. As long as we do not need generalization—which is only necessary when additions are performed within a loop—our analysis (Section 6) is nothing more than a symbolic evaluation, where instead of maintaining a concrete state as a vector of numbers and applying the operations numerically, we maintain AMPs that represent the computation so far, and apply symbolic arithmetics to them. While $\llbracket \mathbb{C} \rrbracket^S$ represents all the possible effects of \mathbb{C} , as a set, we develop in this section a non-deterministic evaluator, which only follows one possible path. Thus a **choose** command will be implemented verbatim—by non-deterministically choosing one branch or the other. Similarly, in a loop, the non-deterministic evaluator literally iterates its body a non-deterministic number of times. The information that it has to maintain, besides the current command, is just the current symbolic state. Importantly, since we are provided with an initial univariate state, all our computation is with such states, which are much more compact than AMPs. Given a degree bound d , a univariate state can be represented by a vector in $[d]^n$. This representation requires polynomial space.

Next, we give a set of definitions for a simple non-deterministic interpreter Ev_S , which does not handle generalization (we give this simple version first for the sake of presentation). Throughout this presentation, \mathbf{s} refers to a univariate state.

The interpretation of **skip**, assignment and multiplication commands is the obvious

$$\text{Ev}_S \llbracket \mathbf{skip} \rrbracket \mathbf{s} \Rightarrow \mathbf{s}; \quad \text{Ev}_S \llbracket \mathbf{X}_i := \mathbf{X}_j \rrbracket \mathbf{s} \Rightarrow \mathbf{set}_{ij} \bullet \mathbf{s}; \quad \text{Ev}_S \llbracket \mathbf{X}_i := \mathbf{X}_j * \mathbf{X}_k \rrbracket \mathbf{s} \Rightarrow \mathbf{mul}_{ijk} \bullet \mathbf{s}$$

where **set**, **mul** are as in the last section, and \Rightarrow is the evaluation relation. Since the addition of abstract univariate monomials x^a and x^b is $x^{\max(a,b)}$, we could simulate addition $\mathbf{X}_i := \mathbf{X}_j + \mathbf{X}_k$ by replacing $\mathbf{s}[i]$ with $\max(\mathbf{s}[j], \mathbf{s}[k])$; however we prefer to use non-determinism and define two alternatives for evaluating this command:

$$\begin{aligned} \text{Ev}_S \llbracket \mathbf{X}_i := \mathbf{X}_j + \mathbf{X}_k \rrbracket \mathbf{s} &\Rightarrow \mathbf{set}_{ij} \bullet \mathbf{s} \\ \text{Ev}_S \llbracket \mathbf{X}_i := \mathbf{X}_j + \mathbf{X}_k \rrbracket \mathbf{s} &\Rightarrow \mathbf{set}_{ik} \bullet \mathbf{s} \end{aligned}$$

This non-determinism means that we can get the precise result as well as another one which is lower and seems redundant. However, this redundancy turns out to be advantageous later, when we get to multivariate bounds.

The interpretation of the choice and sequencing commands is natural, using non-determinism to implement **choose**.

$$\begin{aligned}
\text{EV}_S[\text{choose } C_1 \text{ or } C_2]\mathbf{s} &\Rightarrow \text{EV}_S[C_1]\mathbf{s} \\
\text{EV}_S[\text{choose } C_1 \text{ or } C_2]\mathbf{s} &\Rightarrow \text{EV}_S[C_2]\mathbf{s} \\
\text{EV}_S[C_1; C_2]\mathbf{s} &\Rightarrow \text{EV}_S[C_2](\text{EV}_S[C_1]\mathbf{s})
\end{aligned}$$

For the loop, since we are assuming that there is no addition—hence generalization is not necessary—we just compute composition-closure, i.e., iterate the loop body any finite number of times, using the auxiliary function $\text{EV}_S^*[[C]]$:

$$\begin{aligned}
\text{EV}_S^*[[C]]\mathbf{s} &\Rightarrow \mathbf{s} \\
\text{EV}_S^*[[C]]\mathbf{s} &\Rightarrow \text{EV}_S^*[[C]](\text{EV}_S[[C]]\mathbf{s})
\end{aligned}$$

$$\text{EV}_S[\text{loop } X \{C\}]\mathbf{s} = \text{EV}_S^*[[C]]\mathbf{s}$$

For any \mathbf{x} , the set $[[C]]^S \bullet \mathbf{x}$ and the set of possible results of evaluating $\text{EV}_S[[C]]\mathbf{x}$ have the same maximal elements; thus they define the same worst-case bound.

The difference between computing with full AMPs and evaluating with univariate states as above can be explained as follows: the effect of any computation path of the program is given by an expression

$$\mathbf{p}_1\mathbf{p}_2 \dots \mathbf{p}_j$$

where \mathbf{p}_i represents the i th atomic command (i.e., an assignment) in the path, and juxtaposition $\mathbf{p}\mathbf{q}$ means applying \mathbf{q} after \mathbf{p} (left-to-right composition). In order to evaluate the result given an initial state \mathbf{i} , we compute

$$\mathbf{i}\mathbf{p}_1\mathbf{p}_2 \dots \mathbf{p}_j.$$

The point is that thanks to associativity, we can parenthesize such an expression in different ways. The expression

$$\mathbf{i}(\mathbf{p}_1\mathbf{p}_2 \dots \mathbf{p}_j)$$

represents transforming the path to an AMP (our symbolic semantics from the last section) and then applying the AMP to the given initial state. On the other hand, EV_S implements the computation

$$((\mathbf{i}\mathbf{p}_1)\mathbf{p}_2) \dots \mathbf{p}_j$$

(which is our default reading of a product—left-associative), which gains efficiency, since we only maintain a univariate state.

7.2. Data-flow matrices. When we consider loops that include addition, we have to introduce generalization, in order to account for repeated increments. From the definition of the function LC (Eq. 6.2), we can represent each of the values of $LC([[C]]^S)$ by an expression of this form:

$$\mathbf{p}_{11}\mathbf{p}_{12} \dots \mathbf{p}_{1j_1} (\mathbf{p}_{21}\mathbf{p}_{22} \dots \mathbf{p}_{2j_2})^\tau \mathbf{p}_{31}\mathbf{p}_{32} \dots \mathbf{p}_{3j_3} (\mathbf{p}_{41}\mathbf{p}_{42} \dots \mathbf{p}_{4j_4})^\tau \dots \quad (7.1)$$

where each \mathbf{p}_{ij} is an AMP from $[[C]]^S$, representing a single iteration.

Our interpreter will start with a given univariate state and repeatedly apply the body of the loop; due to non-determinism, each application is equivalent to the application of one possible AMP from $[[C]]^S$, so we are effectively applying $\mathbf{p}_{11}\mathbf{p}_{12} \dots$. However, the parenthesized expressions need a different handling. In order to compute $(\mathbf{p})^\tau$, even if we only want to evaluate it on a univariate state—i.e., $\mathbf{i}(\mathbf{p})^\tau$ —we need to know more about \mathbf{p} :

namely, we have to identify self-dependent variables, and to ensure that generalization is sound (i.e., only applied to iterative fragments). In order to track variable dependences and verify self-dependence, we now introduce *data-flow matrices* and add their maintenance to our interpreter.

Definition 7.1 (Data-Flow Matrix). Given $\mathbf{p} \in \text{AMPo1}$, its data-flow matrix $\text{DFM}(\mathbf{p})$ is a Boolean matrix $M \in M_n(\mathbb{B})$ such that M_{ij} is 1 if $\mathbf{p}[j]$ depends on x_i , and 0 otherwise.

Example 7.2. A data-flow matrix can be associated with any sequence of assignments. Consider the command $X_1 := X_1 * X_2; X_3 := X_2 + X_3$, in a program with $n = 3$ variables, it is represented by the AMP $\mathbf{p} = \langle x_1 x_2, x_2, x_2 + x_3 \rangle$. The data-flow matrix is

$$\text{DFM}(\mathbf{p}) = \begin{bmatrix} 1 & 0 & 0 \\ 1 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

In previous works [NW06, JK09, BJK08], matrices were not Boolean, and the values of their components used to differentiate types of value dependence (the precise meaning of “type” varies). For our algorithm, we find that we only need to track one type of dependence, namely linear dependence. Non-linear dependence will not be tracked at all (it could be, but this would complicate the algorithm; we go by the rule of maintaining the minimal information necessary for a purpose).

Definition 7.3 (Linear Data-Flow Matrix). Given $\mathbf{p} \in \text{AMPo1}$, its *linear* data-flow matrix $\text{LDFM}(\mathbf{p})$ is a Boolean matrix $M \in M_n(\mathbb{B})$ such that M_{ij} is 1 if $\mathbf{p}[j]$ depends on x_i *linearly*, and 0 otherwise.

Example 7.4. For \mathbf{p} as in Example 7.2,

$$\text{LDFM}(\mathbf{p}) = \begin{bmatrix} 0 & 0 & 0 \\ 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}.$$

In order to compute data-flow matrices for composite commands we use ordinary Boolean matrix product. The reader is invited to verify that, if A represents a command C and B represents command D , then $A \cdot B$ represents the data-flow in $C;D$. This works both for DFM and for LDFM . In terms of AMPs, we have

Lemma 7.5. *The matrix abstraction commutes with left-to-right composition: $\text{DFM}(\mathbf{p}\mathbf{q}) = \text{DFM}(\mathbf{p})\text{DFM}(\mathbf{q})$ and $\text{LDFM}(\mathbf{p}\mathbf{q}) = \text{LDFM}(\mathbf{p})\text{LDFM}(\mathbf{q})$.*

Let us define some notations. I is the identity matrix (we always assume that the dimension is fixed by context; it is the number of variables in the program under analysis). The “assignment” matrix $A(j, i)$ represents the data-flow in the command $X_i := X_j$. It differs from I in a single column, column i , which is the j th unit vector:

$$A(j, i) = \begin{matrix} & & & & i \\ & & & & \\ & & & & \\ & & & & \\ & & & & \\ & & & & \\ & & & & \\ j & \left[\begin{array}{ccc} 1 & 0 & \\ & \ddots & 0 \\ & & 0 \\ & & 1 & \ddots \\ & & 0 & & 1 \end{array} \right] \end{matrix}$$

$$\begin{aligned}
\text{EV}_M[\text{skip}]s &\Rightarrow (s, I) \\
\text{EV}_M[X_i := X_j]s &\Rightarrow (\text{set}_{ij} \bullet s, A(j, i)) \\
\text{EV}_M[X_i := X_j * X_k]s &\Rightarrow (\text{mul}_{ijk} \bullet s, E(i)) \\
\text{EV}_M[X_i := X_j + X_k]s &\Rightarrow (\text{set}_{ij} \bullet s, A(j, i) + A(k, i)) \\
\text{EV}_M[X_i := X_j + X_k]s &\Rightarrow (\text{set}_{ik} \bullet s, A(j, i) + A(k, i)) \\
\text{EV}_M[\text{choose } C_1 \text{ or } C_2]s &\Rightarrow \text{EV}_M[C_1]s \\
\text{EV}_M[\text{choose } C_1 \text{ or } C_2]s &\Rightarrow \text{EV}_M[C_2]s \\
\text{EV}_M[C_1; C_2]s &\Rightarrow \text{let } (s_1, M_1) = \text{EV}_M[C_1]s; \text{ let } (s_2, M_2) = \text{EV}_M[C_2]s_1 \\
&\quad \text{in } (s_2, M_1 \cdot M_2)
\end{aligned}$$

Figure 2: The main function of the analyzer, except for loop handling.

We similarly define $E(i)$ (“erase”) to differ from I in having the i th column zero. This represents the LDFM of $X_i := X_j * X_k$. The following observations will be useful.

- Let M be an LDFM representing some transformation \mathbf{p} . Then $I \wedge M$ (the \wedge is applied component-wise) gives the self-dependent variables of \mathbf{p} (as the non-zero diagonal elements).
- $(I \wedge M) \cdot M$ is precisely $\text{LDFM}(\langle \mathbf{p} \rangle)$.

7.3. Maintaining Data-Flow Matrices in Non-Looping Commands. Next, we give to each non-loop command a (non-deterministic) semantics in terms of both a univariate state and a data-flow matrix. The evaluator EV_M computes this semantics. It is invoked as $\text{EV}_M[C]s$, to evaluate command C over initial univariate state s . This returns a pair: the final state and a data-flow matrix, abstracting the computation that took the given state to the final one. See Figure 2. Note that in the case of addition, even if we use the fact that the sum of two abstract monomials equals one of them, we still record the data-flow from both summands.

7.4. Handling loops. When we handle a loop, we will be (in essence) evaluating expressions of the type in Section (7.1). We now outline the ideas used for simulating a parenthesized sub-expression $\langle \mathbf{p}_1 \mathbf{p}_2 \dots \mathbf{p}_j \rangle^\tau$.

We introduce the following notation for changing a set of components of \mathbf{x} :

Definition 7.6. For $Z \subset [n]$, and an n -tuple \mathbf{x} , $\mathbf{x}[Z \leftarrow 0]$ is obtained by changing the components indexed by Z to 0.

We can use this notation with univariate states as well as with (abstract) MPs. In fact, we note that $\mathbf{x}[Z \leftarrow 0] = \text{Id}[Z \leftarrow 0] \bullet \mathbf{x}$.

Let us recall how, in the Closure Algorithm, we arrive at $\langle \mathbf{p} \rangle^\tau$ for some $\mathbf{p} \in \text{Cl}(\llbracket C \rrbracket^S)$.

- (1) We form $\mathbf{p} \in \text{Cl}(\llbracket C \rrbracket^S)$ as $\mathbf{p} = \mathbf{p}_1 \mathbf{p}_2 \dots \mathbf{p}_j$ where every $\mathbf{p}_i \in \llbracket C \rrbracket^S$ represents the effect of an iteration.
- (2) We reduce \mathbf{p} to $\langle \mathbf{p} \rangle$. This involves deleting monomials which are not self-dependent.
- (3) We apply generalization to form $\langle \mathbf{p} \rangle^\tau$.

$$\begin{aligned}
\text{Ev}_M[\llbracket \text{loop } x_\ell \{C\} \rrbracket \mathbf{s}] &\Rightarrow \text{Ev}_M^{\otimes}[\llbracket C \rrbracket](\ell, \mathbf{s}) \\
\text{Ev}_M^*[\llbracket C \rrbracket \mathbf{s}] &\Rightarrow (\mathbf{s}, I) \\
\text{Ev}_M^*[\llbracket C \rrbracket \mathbf{s}] &\Rightarrow \text{let } (\mathbf{s}_1, M_1) = \text{Ev}_M[\llbracket C \rrbracket \mathbf{s}]; \text{ let } (\mathbf{s}_2, M_2) = \text{Ev}_M^*[\llbracket C \rrbracket \mathbf{s}_1] \\
&\quad \text{in } (\mathbf{s}_2, M_1 \cdot M_2) \\
\text{Ev}_M^{\otimes}[\llbracket C \rrbracket](\ell, \mathbf{s}) &\Rightarrow (\mathbf{s}, I) \\
\text{Ev}_M^{\otimes}[\llbracket C \rrbracket](\ell, \mathbf{s}) &\Rightarrow \text{let } Z \subset [n]; \text{let } (\mathbf{s}_1, M_1) = \text{GEN}(Z, \ell, \text{Ev}_M^*[\llbracket C \rrbracket](\mathbf{s}[Z \leftarrow 0])); \\
&\quad \text{let } (\mathbf{s}_2, M_2) = \text{Ev}_M^{\otimes}[\llbracket C \rrbracket](\ell, \mathbf{s}_1) \\
&\quad \text{in } (\mathbf{s}_2, M_1 \cdot M_2)
\end{aligned}$$

Figure 3: Loop handling (up to the *GEN* function). The construct `let $Z \subset [n]$` non-deterministically selects a set.

Our PSPACE algorithm never forms \mathbf{p} . But it goes through $\mathbf{p}_1, \mathbf{p}_2 \dots \mathbf{p}_j$ while simulating their action on an initial univariate state, say \mathbf{x} . If we always do this faithfully, we end up with a result that corresponds to an application of \mathbf{p} to the univariate state, $\mathbf{p} \bullet \mathbf{x}$. Now we have two challenges:

- (1) how to obtain $(\mathbf{p}) \bullet \mathbf{x}$, which may differ from $\mathbf{p} \bullet \mathbf{x}$.
- (2) how to simulate generalization.

The solution to both challenges involves two tools: first, the data-flow matrix; second, we modify the initial state for the simulation of this sequence of loop iterations by non-deterministically setting a set of components to *zero*. Let us explain how this handles the challenges.

- (1) (\mathbf{p}) is obtained from \mathbf{p} by deleting all non-iterative monomials. It is easy to see that

$$(\mathbf{p}) \bullet \mathbf{x} = \mathbf{p} \bullet (\mathbf{x}[\overline{\text{SD}(\mathbf{p})} \leftarrow 0]),$$

where by \overline{S} we mean the complement of S . Our interpreter will guess a set Z and replace \mathbf{x} with $\mathbf{x}[Z \leftarrow 0]$. After simulating a number of iterations of the loop—corresponding to the application of some $\mathbf{p} \in \text{Cl}(\llbracket C \rrbracket^S)$ —we verify that we made a valid guess, meaning that all variables not in Z are self-dependent in \mathbf{p} , using the associated data-flow matrix. Note that a set larger than $\overline{\text{SD}(\mathbf{p})}$ is considered valid; this is important for the next point.

- (2) Consider $\mathbf{p}[i]$ of the form $x_i + q$, where $q(x_1, \dots, x_n)$ depends only on self-dependent variables: such an entry is subject to generalization, multiplying q by τ , which is later replaced by the loop bound x_ℓ . In algorithm *GEN* below, we shall multiply by x_ℓ directly without the intermediate use of the τ symbol. So the expected result is $x_i + x_\ell \cdot q$, which, when evaluated with abstract univariate states, is $\max(x_i, x_\ell \cdot q)$. We see that we can get a correct result in the following way: we non-deterministically either set x_i to 0, and at the end multiply the i 'th component of the result by x_ℓ , or we modify neither x_i nor the final state. In the first case we compute $x_\ell \cdot q$, and in the second $\max(x_i, q)$. The latter result is always sound (it skips generalization), while the former is sound as long as generalization can be rightfully applied, a condition which can be checked using the data-flow matrix (namely that we have evaluated an iterative MP and that $i \in \text{SD}(\mathbf{p})$).

The code for interpreting a loop is presented in Figure 3. The non-deterministic alternatives

for EV_M^* and EV_M^\otimes correspond to zero or more loop iterations respectively. Note how the definition of EV_M^\otimes corresponds to the expression $\text{Cl}(\text{Gen}(\text{Cl}(\llbracket \mathbb{C} \rrbracket^S)))$ in Equation 6.2: an element of the inner closure, $\text{Cl}(\llbracket \mathbb{C} \rrbracket^S)$, corresponds to a simulation of a finite number of iterations, as done by function EV_M^* ; an element of the outer closure is a composition of a finite number of elements, some of which are generalized (we actually generalize whenever possible).

Function GEN checks, using the data-flow matrix, whether we have indeed simulated an iterative MP, and if we did, it applies generalization by multiplying by x_ℓ the outcome of self-dependent variables that have been initialized to zero in this sub-computation. As explained above, the effect is that we multiply the increments (the term q where $\mathbf{p}[i] = x_i + q$).

$$\text{GEN}(Z, \ell, (\mathbf{s}, M)) = \text{if } (\forall i \notin Z : M_{ii} = 1) \text{ then } (\mathbf{s}^{\text{gen}}, M^{\text{gen}}) \text{ else } (\mathbf{s}, M)$$

where

$$\mathbf{s}^{\text{gen}} = \mathbf{s} \prod_{i \in Z : M_{ii} = 1} \mathbf{mul}_{i\ell} \quad M^{\text{gen}} = (I \wedge M)M(I \wedge \overline{M}) \vee (I \wedge M)$$

Note that the multiplication is performed by applying $\mathbf{mul}_{i\ell}$ for each such i . There may be several such indices, hence the use of \prod notation in \mathbf{s}^{gen} (it expresses iteration of postfix composition). The formula for M^{gen} deserves some explanation. Note that it does not depend on Z at all. This is important for the correctness proof. Informally, assuming we have reached the current state by simulating a sequence of assignments equivalent to an AMP \mathbf{p} , the expression $(I \wedge M)M$ is $\text{LDFM}(\llbracket \mathbf{p} \rrbracket)$; multiplying by $(I \wedge \overline{M})$ results in *removing* from this matrix the columns that correspond to self-dependent entries; and then we put them back by adding $I \wedge M$. The point of this exercise can be seen in the next example.

Example 7.7. We illustrate an application of function GEN . We refer back to Example 6.1. Recall that the analysis of the loop body produced the following AMPs

$$\mathcal{S} = \{ \langle x_1 + x_3, x_2 + x_4, x_3, x_3 \rangle, \quad \langle x_1, x_2 + x_3 + x_4, x_3, x_3 \rangle \}.$$

We focus on the path corresponding to the first AMP, and consider its evaluation with initial state $\mathbf{s} = \langle x^1, x^2, x^3, x^4 \rangle$. Applying EV_M^* , we may obtain the final result $\mathbf{s}' = \langle x^3, x^4, x^3, x^3 \rangle$ along with the matrix

$$M = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 0 & 0 \end{bmatrix}.$$

What GEN does depends on the choice of Z . In the case $Z = \emptyset$, it returns the above result, which is equivalent to skipping generalization. Instead, EV_M^\otimes may choose a non-empty set; for example, let $Z = \{1, 4\}$. Then EV_M^\otimes modifies \mathbf{s} by setting the first and fourth component to zero before invoking EV_M^* ; so its result changes into $\mathbf{s}'' = \langle x^3, x^2, x^3, x^3 \rangle$, while the matrix is unaffected. Next, function GEN checks the matrix: it confirms the condition for generalization, as $M_{22} = M_{33} = 1$. Only $i = 1$ satisfies $i \in Z \wedge M_{ii} = 1$, so just this component is multiplied by the loop bound $\mathbf{s}''[3] = x^3$ and we get $\mathbf{s}^{\text{gen}} = \langle x^6, x^2, x^3, x^3 \rangle$. For clarity, we now break the computation of M^{gen} in parts: first,

$$M_1 = (I \wedge M)M = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

removes the dependence on x_4 , which is not self-dependent. In other words, M_1 represents the iterative part of a corresponding AMP. Then,

$$M_1(I \wedge \overline{M}) \vee (I \wedge M) = \begin{bmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix} \vee \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 \end{bmatrix}$$

effectively replaces the first column with a unit column. Why? In M , we had two 1's in this column, corresponding to the expression $x_1 + x_3$. Generalization turns this into $x_1 + x_3 \cdot x_3$. But we are only interested in linear terms, so we replace this column with a unit column, erasing the dependence on x_3 .

8. COMPLEXITY

We now consider the space complexity of the algorithm. The size of a univariate state, implemented as a degree vector, is $O(n(1 + \log d))$ bits where $d \geq 1$ is the highest degree reached. In the problem formulation where a degree is specified as input, the algorithm can safely truncate any higher degree down to d , as this will still give a correct answer to the user's question. It remains to consider the space occupied by the recursion stack. This will be proportional to the depth of the program's syntax tree once we avoid the calls of functions EV_M^* and EV_M^\otimes to themselves, or rather change them into tail calls. This can be done by routine rewriting of the functions, using an auxiliary "accumulator" parameter, so for instance the code

$$\begin{aligned} \text{EV}_M^*[\mathbf{C}]\mathbf{s} &\Rightarrow (\mathbf{s}, I) \\ \text{EV}_M^*[\mathbf{C}]\mathbf{s} &\Rightarrow \text{let } (\mathbf{s}_1, M_1) = \text{EV}_M[\mathbf{C}]\mathbf{s}; \quad \text{let } (\mathbf{s}_2, M_2) = \text{EV}_M^*[\mathbf{C}]\mathbf{s}_1 \\ &\quad \text{in } (\mathbf{s}_2, M_1 \cdot M_2) \end{aligned}$$

changes into

$$\begin{aligned} \text{TAILEV}_M^*[\mathbf{C}](\mathbf{s}, A) &\Rightarrow (\mathbf{s}, A) \\ \text{TAILEV}_M^*[\mathbf{C}](\mathbf{s}, A) &\Rightarrow \text{let } (\mathbf{s}_1, M_1) = \text{EV}_M[\mathbf{C}]\mathbf{s} \\ &\quad \text{in TAILEV}_M^*[\mathbf{C}](\mathbf{s}_1, A \cdot M_1) \end{aligned}$$

and calls to EV_M^* from EV_M^\otimes pass I for the accumulator parameter. Thus we have

Theorem 8.1. *The univariate polynomial-bound problem (Problem 3.5) as well as the lower-bound generation problem (Problem 3.6) has polynomial space complexity.*

Note that the non-deterministic nature of the algorithm places the decision problem in NPSpace, however by Savitch's theorem this implies PSPACE as well. The corresponding generation problem can also be determinized by essentially turning it into an exhaustive search, still in polynomial space. Finally we would like to recall that a program in our language can have exponentially-growing variables as well, and emphasize that handling them, as described in [BAH20], does not increase the complexity of our problem (the output in such a case will simply indicate that the variable has an exponential lower bound; no tight bound will be given).

9. PROOF OF THE PSPACE ALGORITHM

The purpose of this section is to show that our PSPACE algorithm, EV_M , obtains correct results. How do we define correctness? The original goal of the algorithm is to obtain symbolic expressions that tightly bound the concrete *numeric* results obtained by a core-language program, applied to an initial integer-valued state. However, thanks to our previous work, we know that such a set of symbolic expressions—namely AMPs—is obtained by the abstract semantics of Section 6. So now we are able to define $\llbracket \mathcal{C} \rrbracket^S$ as our reference and define our goal as matching the results it provides, specialized to a univariate initial state. So the purpose of this section is to prove the following:

Theorem 9.1. *The interpreter EV_M satisfies these two correctness claims:*

- (1) *Soundness: Given any univariate state \mathbf{x} , if $EV_M \llbracket \mathcal{C} \rrbracket \mathbf{x} \Rightarrow^* (\mathbf{y}, M)$ (for any M) then $\exists \mathbf{p} \in \llbracket \mathcal{C} \rrbracket^S$ such that $\mathbf{y} \leq \mathbf{p} \bullet \mathbf{x}$.*
- (2) *Completeness: Given any $\mathbf{p} \in \llbracket \mathcal{C} \rrbracket^S$, and univariate state \mathbf{x} , there is $\mathbf{y} \geq \mathbf{p} \bullet \mathbf{x}$ such that $EV_M \llbracket \mathcal{C} \rrbracket \mathbf{x} \Rightarrow^* (\mathbf{y}, M)$ for some M .*

Note that we can also state this as follows: for any \mathbf{x} , the set $\llbracket \mathcal{C} \rrbracket^S \bullet \mathbf{x}$ and the set of possible results of evaluating $EV_M \llbracket \mathcal{C} \rrbracket \mathbf{x}$ have the same maximal elements; thus they define the same worst-case bound.

Obviously, the proof is inductive and to carry it out completely we also have to assert something about the matrices. We define a notation for brevity:

Definition 9.2. Let f be a program function that, given a univariate state, evaluates to a pair—a univariate state and a matrix. Let \mathcal{S} be a set of AMPs. We write:

- $f \hookrightarrow \mathcal{S}$ if $\forall \mathbf{x}$, if $f \mathbf{x} \Rightarrow^* (\mathbf{y}, M)$ then $\exists \mathbf{p} \in \mathcal{S}$ such that $\mathbf{y} \leq \mathbf{p} \bullet \mathbf{x}$ and $M = \text{LDFM}(\mathbf{p})$.
- $f \leftarrow \mathcal{S}$ if $\forall \mathbf{p} \in \mathcal{S}$, $\forall \mathbf{x}$, $\exists (\mathbf{y}, M)$ such that $f \mathbf{x} \Rightarrow^* (\mathbf{y}, M)$, $\mathbf{y} \geq \mathbf{p} \bullet \mathbf{x}$ and $M = \text{LDFM}(\mathbf{p})$.

The full correctness claim will be the following:

Theorem 9.3. *The interpreter EV_M satisfies these two correctness claims:*

- (1) *Soundness: Given any command \mathcal{C} , $EV_M \llbracket \mathcal{C} \rrbracket \hookrightarrow \llbracket \mathcal{C} \rrbracket^S$.*
- (2) *Completeness: Given any command \mathcal{C} , $EV_M \llbracket \mathcal{C} \rrbracket \leftarrow \llbracket \mathcal{C} \rrbracket^S$.*

Correctness is straight-forward for straight-line code (this basically amounts to the associativity argument proposed in Section 7.1, with Lemma 7.5 justifying the matrices), and it trivially extends to commands with branching, and even to loops without additions (since they are analyzed by unrolling finitely many times). It is where *generalization* is used that correctness is subtle and so, this is the focus of this section. Since the case of straight-line (or more generally, loop-free) code is simple, we will state the properties without a detailed proof.

9.1. Correctness without loops.

Lemma 9.4. *Let \mathcal{C} be a loop-free command. Then $EV_M \llbracket \mathcal{C} \rrbracket \hookrightarrow \llbracket \mathcal{C} \rrbracket^S$ and $EV_M \llbracket \mathcal{C} \rrbracket \leftarrow \llbracket \mathcal{C} \rrbracket^S$.*

EV_M^* just iterates EV_M an undetermined number of times, which matches the definition of composition-closure precisely so the following is also straight-forward.

Lemma 9.5. *Let \mathcal{C} be a loop-free command. Then $EV_M^* \llbracket \mathcal{C} \rrbracket \hookrightarrow Cl(\llbracket \mathcal{C} \rrbracket^S)$ and $EV_M^* \llbracket \mathcal{C} \rrbracket \leftarrow Cl(\llbracket \mathcal{C} \rrbracket^S)$.*

9.2. Soundness of the analysis of loops. We now wish to extend the soundness claim to loops. The proof is by structural induction, where all loop-free commands are covered by the above lemmas, and similarly, given correctness for commands \mathbf{C} and \mathbf{D} , correctness for their sequential composition, or non-deterministic choice, follow straight from definitions. So the main task is to prove the correctness for loops, namely the correctness of Ev_M^* . The main task is to relate the possible results of $\text{GEN}(Z, \ell, \text{Ev}_M^*[\mathbf{C}](\mathbf{x}[Z \leftarrow 0]))$, where Z ranges over subsets of $[n]$, with members of the AMP set $\text{Gen}(Cl(\llbracket \mathbf{C} \rrbracket^S))[x_\ell/\tau]$. Note that in the expression $\text{Gen}(Cl(\llbracket \mathbf{C} \rrbracket^S))[x_\ell/\tau]$, unlike in Section 6, we substitute the x_ℓ for τ immediately after generalization; this is done for the sake of comparing the results, since Ev_M does not use τ . This modification of the Closure Algorithm is harmless, as for AMPs \mathbf{p}, \mathbf{q} from this set, we can use the rule $(\mathbf{p} \bullet \mathbf{q})[x_\ell/\tau] = (\mathbf{p}[x_\ell/\tau]) \bullet (\mathbf{q}[x_\ell/\tau])$, due to the assumption that \mathbf{X}_ℓ is not modified inside the loop.

Lemma 9.6. *For any univariate state \mathbf{x} , and $Z \subset [n]$, if*

$$\text{GEN}(Z, \ell, \text{Ev}_M^*[\mathbf{C}](\mathbf{x}[Z \leftarrow 0])) \Rightarrow^* (\mathbf{y}, M),$$

then there is $\mathbf{r} \in \text{Gen}(Cl(\llbracket \mathbf{C} \rrbracket^S))$ such that $\mathbf{y} \leq \mathbf{r}[x_\ell/\tau] \bullet \mathbf{x}$, and $M = \text{LDFM}(\mathbf{r}[x_\ell/\tau])$.

Note that this lemma constitutes an induction step; it is used under the inductive assumption that for the loop body \mathbf{C} , we have soundness.

Proof. Let $\text{Ev}_M^*[\mathbf{C}](\mathbf{x}[Z \leftarrow 0]) \Rightarrow^* (\mathbf{y}, M)$. When generalization is not applied, i.e., $\mathbf{y} = \mathbf{s}'$ and $M = M'$, we have, by the inductive assumption, $\mathbf{p} \in Cl(\llbracket \mathbf{C} \rrbracket^S)$ such that $\mathbf{s}' \leq \mathbf{p} \bullet (\mathbf{x}[Z \leftarrow 0]) \leq \mathbf{p} \bullet \mathbf{x}$ and $M' = \text{LDFM}(\mathbf{p})$. Further, $\mathbf{p} = \mathbf{p}[x_\ell/\tau]$ (since there is no τ).

Next, suppose that matrix M' satisfies the condition for generalization ($\forall i \notin Z : M'_{ii} = 1$). By the inductive assumption, we can choose $\mathbf{p} \in Cl(\llbracket \mathbf{C} \rrbracket^S)$ such that $\mathbf{s}' \leq \mathbf{p} \bullet \mathbf{x}[Z \leftarrow 0]$ and $M' = \text{LDFM}(\mathbf{p})$. We now wish to show that $\mathbf{y} = \mathbf{s}^{\text{gen}} \leq (\mathbf{p})^\tau[x_\ell/\tau] \bullet \mathbf{x}$. We focus on the non-trivial case which is that of an index k to which generalization applies, i.e., $k \in Z$ and $M'_{kk} = 1$. Thus x_k is self-dependent in \mathbf{p} . Let $\mathbf{p}_Z = \mathbf{p} \bullet \text{Id}[Z \leftarrow 0]$, then \mathbf{p}_Z is iterative. In fact, $\mathbf{p}_Z \sqsubseteq (\mathbf{p})$. Write $(\mathbf{p})[k] = x_k + q_k$, where q_k is a polynomial not dependent on x_k . Since $k \in Z$, $\mathbf{p}_Z[k] \sqsubseteq q_k$. Moreover, $\mathbf{s}'[k] \leq \mathbf{p}[k] \bullet \mathbf{x}[Z \leftarrow 0] = \mathbf{p}_Z[k] \bullet \mathbf{x}$. Now, in the Closure Algorithm, function Gen sets $(\mathbf{p})^\tau[k] = x_k + q_k\tau$ (the τ subsequently substituted with x_ℓ), so multiplying $\mathbf{s}'[k]$ by $\mathbf{s}'[\ell]$, as GEN does, produces

$$\mathbf{s}^{\text{gen}}[k] \leq (\mathbf{p}_Z[k] \bullet \mathbf{x}) \cdot \mathbf{s}'[\ell] \leq ((\mathbf{p})^\tau[k][x_\ell/\tau]) \bullet \mathbf{x}.$$

We conclude that the lemma is fulfilled by $\mathbf{r} = (\mathbf{p})^\tau$.

Regarding the computation of the matrix M^{gen} , we refer to the text in Section 7.4. It explains why, given $M' = \text{LDFM}(\mathbf{p})$, we get $M^{\text{gen}} = \text{LDFM}((\mathbf{p})^\tau[x_\ell/\tau])$. \square

Ev_M^* combines an undetermined number of calls to $\text{GEN}(Z, \ell, \text{Ev}_M^*[\mathbf{C}](\mathbf{x}[Z \leftarrow 0]))$, for arbitrary sets Z , which by straight-forward induction, and the definition of $\llbracket \text{loop } \mathbf{X}_\ell \{ \mathbf{C} \} \rrbracket^S$, gives

Lemma 9.7 (Soundness for loops). *Let \mathbf{C} be a command such that $\text{Ev}_M[\mathbf{C}] \leftrightarrow \llbracket \mathbf{C} \rrbracket^S$. Then*

$$\lambda \mathbf{s}. \text{Ev}_M^*[\mathbf{C}](\ell, \mathbf{s}) \leftrightarrow Cl(\text{Gen}(Cl(\llbracket \mathbf{C} \rrbracket^S)))[x_\ell/\tau].$$

Hence,

$$\text{Ev}_M[\llbracket \text{loop } \mathbf{X}_\ell \{ \mathbf{C} \} \rrbracket] \leftrightarrow \llbracket \text{loop } \mathbf{X}_\ell \{ \mathbf{C} \} \rrbracket^S.$$

9.3. Completeness for an unnested loop. Again, the crux of this proof will be the application of generalization. Here we want to show that when the Closure Algorithm generalizes, EV_M^* can match its result. Based on the presentation in Section 7.4, we might expect that for $\mathbf{p} \in \text{Cl}(\llbracket \mathbb{C} \rrbracket^S)$, given that EV_M^* can match $\mathbf{p} \bullet \mathbf{x}$, we will capture the results of generalization by making a correct guess of the self-dependent variables and the accumulators. Let us first formalize the latter term.

Definition 9.8. Let $\mathbf{p} \in \text{AMPo1}$ be iterative. Recall that for $i \in \text{SD}(\mathbf{p})$ where $\mathbf{p}[i] = x_i + q_i$,

$$\mathbf{p}^\tau[i] = x_i + q_i \cdot \tau$$

So that, given a univariate state \mathbf{x} ,

$$\mathbf{p}^\tau[i][x_\ell/\tau] \bullet \mathbf{x} = \max((q_i \bullet \mathbf{x}) \cdot \mathbf{x}[\ell], \mathbf{p}[i] \bullet \mathbf{x}).$$

Let $\text{Acc}(\ell, \mathbf{p}, \mathbf{x})$ be the set of indices i where the first term under the \mathbf{max} is larger; intuitively, where generalization increases the result. We refer to $\text{Acc}(\ell, \mathbf{p}, \mathbf{x})$ as the *accumulators* in the computation under consideration (i.e., iterating \mathbf{p} starting at state \mathbf{x}).

Recall that function GEN is applied to a state-matrix pair (\mathbf{s}', M) obtained from $\text{EV}_M^* \llbracket \mathbb{C} \rrbracket \mathbf{s}$. By correctness of this analysis (i.e., the inductive assumption), we assume $\mathbf{s}' \geq \mathbf{p} \bullet \mathbf{s}$, where $\mathbf{p} \in \text{Cl}(\llbracket \mathbb{C} \rrbracket^S)$, and $M = \text{LDFM}(\mathbf{p})$. We shall say that this application of EV_M^* simulates \mathbf{p} (in the completeness statement, we quantify over \mathbf{p} first, so it may be considered fixed throughout the discussion). When EV_M^* is applied to a state in which some entries are set to 0, namely, $\mathbf{s}[Z \leftarrow 0]$, this is equivalent to simulating $\mathbf{p} \bullet \text{Id}[Z \leftarrow 0]$ on \mathbf{s} . Then, GEN possibly multiplies some entries of \mathbf{s}' by $\mathbf{s}'[\ell] = \mathbf{s}[\ell]$. This can be seen as modifying $\mathbf{p} \bullet \text{Id}[Z \leftarrow 0]$ into a new AMP \mathbf{q} by multiplying some entries by x_ℓ , before applying the AMP to \mathbf{s} . We express this by writing $\mathbf{q} = \overline{\text{GEN}}(Z, \mathbf{p}, \ell)$, noting that the action turning \mathbf{p} into \mathbf{q} depends on \mathbf{p} , M (which is determined by \mathbf{p}) and Z , but not on \mathbf{s} . We also note that M^{gen} depends only on M .

Based on the presentation of the algorithm, the reader may infer that we intend that, for an iterative \mathbf{p} , letting $A = \text{Acc}(\ell, \mathbf{p}, \mathbf{x})$, we should have $\overline{\text{GEN}}(A, \mathbf{p}, \ell) \bullet \mathbf{x} = \mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x}$. However, this is not always the case, as shown by the following example.

Example 9.9. Let $\mathbf{p} = \langle x_1, x_2 + x_1, x_3 + x_2, x_4 \rangle$, and $\ell = 4$. Note that \mathbf{p} is iterative. Let $\mathbf{x} = \langle x^2, x^2, x^1, x^1 \rangle$. Then $\text{Acc}(\ell, \mathbf{p}, \mathbf{x}) = \{2, 3\}$. Letting $A = \{2, 3\}$, it is easy to check that $\mathbf{q} = \overline{\text{GEN}}(A, \mathbf{p}, \ell) = \langle x_1, x_1 x_4, 0, x_4 \rangle$. Since we have begun our computation with $\mathbf{x}[A \leftarrow 0] = \langle x^2, 0, 0, x^1 \rangle$, the final state is $\langle x^2, x^3, 0, x^1 \rangle$. Note that the result falls short of $\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x} = \langle x^2, x^3, x^3, x^1 \rangle$.

The solution to this mismatch is given by Lemma 9.12 below, which shows that by calling GEN *multiple times* in succession, in other words, by going through a number of recursive calls of EV_M^* , completeness is recovered.

We shall now make some preparations for this lemma.

The Dependence Graph. The matrix $\text{DFM}(\mathbf{p})$ may be seen as the adjacency matrix of a graph, which we call $G(\mathbf{p})$. Arcs in this graph represent data-flow in \mathbf{p} (we use $\text{DFM}(\mathbf{p})$ and not $\text{LDFM}(\mathbf{p})$, so this includes non-linear dependence). Paths in the graph correspond to data-flow effected by iterating \mathbf{p} . For instance, if we have a path $i \rightarrow j \rightarrow k$ in $G(\mathbf{p})$, then x_i will appear in the expression $(\mathbf{p} \circ \mathbf{p})[k]$.

Lemma 9.10. *If \mathbf{p} is iterative, and causes no exponential growth when iterated on concrete data, then $G(\mathbf{p})$ has no cycles of length greater than 1.*

Note that the assumption refers to concrete computation ($\mathbf{p} \circ \mathbf{p} \circ \dots$), and that we are building on the assumption that the loop under analysis is polynomially bounded (this was presumably verified beforehand, see Section 3).

Proof. Assume, to the contrary, that $G(\mathbf{p})$ does have a cycle $i \rightarrow \dots \rightarrow k \rightarrow i$, where $k \neq i$. Let r be the length of the cycle. Since \mathbf{p} is iterative, all these variables must be self-dependent (as some other entry depends on them). Thus we have $\mathbf{p}[i] = x_i + q(x_1, \dots, x_n)$ where q involves x_k ; but $\mathbf{p}[k]$ depends on x_{k-1} , and so on; hence $\mathbf{p}^{(r-1)}[k] \geq x_i$, and $\mathbf{p}^{(r)}[i] \geq 2x_i$. Thus iteration of \mathbf{p} generates exponential growth, contradicting the assumption. \square

When focusing on a particular AMP we may assume, w.l.o.g., that the variables are indexed in an order consistent with $G(\mathbf{p})$, so that if x_i depends on x_j then $j \leq i$. We shall refer to an (abstract) MP satisfying this property as *neat*³.

We also state a rather evident property of an iterative MP:

Lemma 9.11. *If \mathbf{p} is iterative, then $\mathbf{p} \bullet \mathbf{p} \geq \mathbf{p}$. A fortiori, $\mathbf{p} \bullet \mathbf{p}^\tau \geq \mathbf{p}$.*

Proof. Consider the calculation of $(\mathbf{p} \bullet \mathbf{p})[k]$; it is obtained by substituting $\mathbf{p}[i]$ for each occurrence of x_i in $\mathbf{p}[k]$. By iterativity, we know that $\mathbf{p}[i]$ includes x_i , so we can clearly only enlarge the polynomial. \square

Lemma 9.12. *Let $\mathbf{p} \in \text{AMPo1}$ be neat and iterative, and let \mathbf{x} be a univariate state. There are sets Z_1, \dots, Z_{n-1} such that*

$$\mathbf{q}_{n-1} \bullet \mathbf{q}_{n-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x} \geq \mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x}$$

where $\mathbf{q}_i = \overline{\text{GEN}}(Z_i, \mathbf{p}, \ell)$.

Proof. Let $A = \text{Acc}(\ell, \mathbf{p}, \mathbf{x})$. We define sets Z_i for $i = 1, 2, \dots, n-1$ as follows:

$$Z_i = A \cap \{i+1\} \cup \overline{\text{SD}(\mathbf{p})}.$$

We now state an inductive claim and prove it. The claim is: for $1 \leq i \leq n$, for all $j \leq i$,

$$(\mathbf{q}_{i-1} \bullet \mathbf{q}_{i-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x})[j] \geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})[j]. \quad (9.1)$$

Note that (9.1) specializes to the lemma's statement for $i = n$.

To start the induction, consider $i = 1$. Then we only have to prove that

$$\mathbf{x}[1] \geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})[1].$$

By neatness, $\mathbf{p}[1]$ can only depend on x_1 , in fact $\mathbf{p}[1] = x_1$. Also $\mathbf{p}^\tau[1] = x_1$. So equality holds.

Now, let $k > 1$, and assume that (9.1) holds when $i = k-1$. To prove it for $i = k$, we consider some sub-cases.

(1) $j \leq k$ and $j \notin \text{SD}(\mathbf{p})$: in this case, $\mathbf{p}[j]$ only depends on x_t with $t < k$ where, by IH,

$$(\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x})[t] \geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})[t].$$

³Readers who like linear algebra may draw some intuition about neat MPs from thinking about triangular matrices whose diagonal elements are in $\{0, 1\}$. Interestingly, this structure has also popped up in other works in the area [HFG20].

Since $j \notin \text{SD}(\mathbf{p})$, the construction of $\overline{GEN}(Z_{k-1}, \mathbf{p}, \ell)$ does not modify $\mathbf{p}[j]$, i.e.,

$$\begin{aligned} \mathbf{q}_{k-1}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) &= \mathbf{p}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \\ &\geq \mathbf{p}[j] \bullet (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x}) \\ &\geq \mathbf{p}[j] \bullet \mathbf{x} && \text{by Lemma 9.11} \\ &= \mathbf{p}[j]^\tau[x_\ell/\tau] \bullet \mathbf{x} && \text{since } j \notin \text{SD}(\mathbf{p}). \end{aligned}$$

(2) $j < k$ where $j \in \text{SD}(\mathbf{p})$: here $\mathbf{p}[j]$ depends on x_j . Yet $j \notin Z_{k-1}$, so the construction of $\overline{GEN}(Z_{k-1}, \mathbf{p}, \ell)$ does not modify $\mathbf{p}[j]$, i.e.,

$$\begin{aligned} \mathbf{q}_{k-1}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) &= \mathbf{p}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \\ &\geq (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x})[j] && \text{(since } \mathbf{p}[j] \geq x_j) \\ &\geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})[j] && \text{by IH.} \end{aligned}$$

(3) $j = k \in \text{SD}(\mathbf{p})$ but $k \notin A_{k-1}$: here $\mathbf{p}[j]$ depends on x_j . Yet $j \notin Z_{k-1}$, so the construction of $\overline{GEN}(Z_{k-1}, \mathbf{p}, \ell)$ does not modify $\mathbf{p}[j]$, i.e.,

$$\begin{aligned} \mathbf{q}_{k-1}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) &= \mathbf{p}[j] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \\ &\geq \mathbf{p}[j] \bullet \mathbf{x} && \text{(since } \mathbf{q}_i[j] \geq x_j \text{ for all } i) \\ &\geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})[j] && \text{since } j \notin A_{k-1}. \end{aligned}$$

(4) $j = k \in A_{k-1}$: here, $\mathbf{p}[j] = \mathbf{p}[k]$ depends on x_k . We may write $\mathbf{p}[k] = x_k + z$ for a polynomial z in x_1, \dots, x_{k-1} . We have, by IH,

$$z \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \geq z \bullet (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x})$$

Since $k \in A_{k-1}$, we know that

$$\mathbf{p}^\tau[k][x_\ell/\tau] \bullet \mathbf{x} = (z \bullet \mathbf{x}) \cdot \mathbf{x}[k] = (z \cdot x_\ell) \bullet \mathbf{x}$$

Also, since $k \in Z_{k-1}$,

$$\mathbf{q}_{k-1}[k] = \overline{GEN}(Z_{k-1}, \mathbf{p}, \ell)[k] = (\mathbf{p}[k] \bullet \text{Id}[Z_{k-1} \leftarrow 0]) \cdot x_\ell = z \cdot x_\ell$$

so

$$\begin{aligned} \mathbf{q}_{k-1}[k] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) &= (z \cdot x_\ell) \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \\ &= \mathbf{p}^\tau[k][x_\ell/\tau] \bullet (\mathbf{q}_{k-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x}) \\ &\geq \mathbf{p}^\tau[k][x_\ell/\tau] \bullet (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x}) && \text{by IH} \\ &\geq (\mathbf{p}^\tau[x_\ell/\tau] \bullet \mathbf{x}) && \text{by Lemma 9.11.} \quad \square \end{aligned}$$

We are now ready to conclude

Lemma 9.13. *For any univariate state \mathbf{x} , for any $\mathbf{r} \in (\text{Gen}(\text{Cl}(\llbracket \mathbb{C} \rrbracket^S)))$, there are \mathbf{y}, M such that $\text{Ev}_M^{\otimes} \llbracket \mathbb{C} \rrbracket(\ell, \mathbf{x}) \Rightarrow^* (\mathbf{y}, M)$, $\mathbf{y} \geq \mathbf{r}[x_\ell/\tau] \bullet \mathbf{x}$, and $M = \text{LDFM}(\mathbf{r}[x_\ell/\tau])$.*

Proof. We focus on the interesting case, which is when generalization is applied to (\mathbf{p}) , for $\mathbf{p} \in \text{Cl}(\llbracket \mathbb{C} \rrbracket^S)$, to produce \mathbf{r} . Note that $(\mathbf{p}) = \mathbf{p} \bullet \text{Id}[\overline{\text{SD}(\mathbf{p})} \leftarrow 0]$. For any set $Z_i \supseteq \overline{\text{SD}(\mathbf{p})}$, it is easy to see that

$$\mathbf{q}_i \stackrel{\text{def}}{=} \overline{GEN}(Z_i, \mathbf{p}, \ell) = \overline{GEN}(Z_i, (\mathbf{p}), \ell).$$

Using Lemma 9.12 we have sets Z_1, \dots, Z_{n-1} , containing $\overline{\text{SD}(\mathbf{p})}$, such that

$$\mathbf{q}_{n-1} \bullet \mathbf{q}_{n-2} \bullet \dots \bullet \mathbf{q}_1 \bullet \mathbf{x} \geq (\mathbf{p})^\tau[x_\ell/\tau] \bullet \mathbf{x} = \mathbf{r}[x_\ell/\tau] \bullet \mathbf{x}.$$

We conclude, using Lemma 9.5, that starting with \mathbf{x} , and making at most n recursive calls, Ev_M^{\otimes} can reach $\mathbf{y} \geq \mathbf{r}[x_\ell/\tau] \bullet \mathbf{x}$ as desired. The correctness of the matrix is again the equation $M^{\text{gen}} = \text{LDFM}((\mathbf{p})^\tau[x_\ell/\tau])$, which has been argued in Section 7.4. \square

Straight-forward induction, and the definition of $\llbracket \text{loop } X_\ell \{C\} \rrbracket^S$, give

Lemma 9.14 (Completeness for loops). *Let C be a command such that $\text{Ev}_M\llbracket C \rrbracket \leftrightarrow \llbracket C \rrbracket^S$. Then*

$$\lambda \mathbf{s}. \text{Ev}_M^{\otimes}\llbracket C \rrbracket(\ell, \mathbf{s}) \leftrightarrow \text{Cl}(\text{Gen}(\text{Cl}(\llbracket C \rrbracket^S)))[x_\ell/\tau].$$

Hence,

$$\text{Ev}_M\llbracket \text{loop } X_\ell \{C\} \rrbracket \leftrightarrow \llbracket \text{loop } X_\ell \{C\} \rrbracket^S.$$

10. PSPACE-COMPLETENESS OF BOUND ANALYSIS

In this section we complement our PSPACE upper bound with a hardness result, to show that our classification of the problem's complexity as PSPACE is tight. The hardness proof is a reduction from *termination of Boolean Programs*, a known PSPACE-complete problem. First, we state the definition of the decision problem to which we reduce. This is a special case of the univariate-bound decision problem, with a fixed initial state.

Definition 10.1. The decision problem DEG (for “degree”) is defined as the set of triples (P, j, d) such that P is a core-language program, where the maximal value of X_j at the completion of the program, in terms of the univariate input $\mathbf{x} = \langle x, x, \dots, x \rangle$, has a polynomial lower bound of $\Omega(x^d)$.

The complexity of the problem is classified in relation to the “input size” defined as: $|P| + d$, where $|P|$ is the size of the syntax tree representing P . This means that we allow d to be represented in unary notation (and we will find that this does not affect the complexity class).

Definition 10.2. A *Boolean program* is an instruction sequence $b = 1:\text{I}_1 \ 2:\text{I}_2 \dots m:\text{I}_m$ specifying a computation on Boolean variables B_1, \dots, B_k , with program locations $1, 2, \dots, m$.

Instructions I_ℓ have two forms: $X_i := \text{not } X_i$, and *if* X_i *then goto* ℓ' *else* ℓ'' . Here $1 \leq i \leq k$ and $\ell, \ell', \ell'' \in \{0, 1, 2, \dots, m\}$, where ℓ, ℓ', ℓ'' are always three different locations.

Semantics: the *computation* by b is a finite or infinite state sequence $b \vdash (\ell_1, \sigma_1) \rightarrow (\ell_2, \sigma_2) \rightarrow \dots$, where each store σ assigns a truth value in $\{\text{true}, \text{false}\}$ to each of b 's variables, and ℓ_t is the program location at time t .

We are considering *input free* programs. These programs have a fixed initial state: $\ell_1 = 1$, and σ_1 assigns *false* to every variable. Given state (ℓ_t, σ_t) , if $\ell_t = 0$ then the computation has terminated, else the following rules apply.

If instruction I_{ℓ_t} is $X_i := \text{not } X_i$, then σ_{t+1} is identical to σ_t except that $\sigma_{t+1}(B_i) = \neg\sigma_t(B_i)$. Further, $\ell_{t+1} = (\ell_t + 1) \bmod (m + 1)$.

If instruction I_{ℓ_t} is *if* X_i *then goto* ℓ' *else* ℓ'' , then σ_{t+1} is identical to σ_t . Further, $\ell_{t+1} = \ell'$ if $\sigma_t(B_i) = \text{true}$, and $\ell_{t+1} = \ell''$ if $\sigma_t(B_i) = \text{false}$.

Finally, program b *terminates* if for some t : $b \vdash (\ell_1, \sigma_1) \rightarrow \dots \rightarrow (\ell_t, \sigma_t) = (0, \sigma_t)$.

The following lemma is proved in [Jon97, Chapter 28] (with a trivial difference, concerning acceptance instead of termination):

Lemma 10.3. *The following set is complete for PSPACE:*

$$\mathcal{B} = \{b \mid b \text{ is an input-free Boolean program and } b \text{ terminates}\}.$$

Theorem 10.4. *The problem DEG is PSPACE-hard, even for programs with a single loop.*

Proof. Reduction from problem \mathcal{B} above.

Suppose program $b = 1:I_1 \ 2:I_2 \dots m:I_m$ has k variables B_1, \dots, B_k . Without loss of generality, each variable has value *false* after execution, if b terminates (just add at the end of b one test and one assignment for each variable.)

Program p will have $2+2k(m+1)$ variables named X_1, X_2 , and then $X_{\ell,i,v}$ for all $0 \leq \ell \leq m$, $1 \leq i \leq k$ and $v \in \{0, 1\}$. Informally, the program simulates the Boolean program, such that the pair $X_{\ell,i,0}$ and $X_{\ell,i,1}$ represent the value b of variable B_i when program location ℓ was last visited. This is not a deterministic simulation, due to the absence of deterministic branching; instead, the program may take many paths, most of which are wrong simulations; and the contents of variables $X_{\ell,i,v}$ will reflect the correctness of the simulation. In fact, the value of each such variable will be either x or x^2 , where the value x denotes error. That is, if the program reaches location ℓ with $B_i = v$, only the path that simulates it correctly will have $X_{\ell,i,v} = x^2$.

We first define the initialization command INIT: for all i , it sets $X_{1,i,0}$ to X_1^2 (the initial value of B_i is *false*). All other simulation variables are set to X_1 .

For every program location ℓ we define a command C_ℓ “simulating” the corresponding instruction, as follows:

- For instruction $\ell : X_i := \text{not } X_i$,
 C_ℓ modifies only variables $X_{\ell',i,v}$ where $\ell' = (\ell + 1) \bmod (m + 1)$. For each $v = 0, 1$, it sets:

$$X_{\ell',i,v} := X_{\ell,i,\neg v}$$
 while for all $j \neq i$, and $v = 0, 1$, it sets

$$X_{\ell',i,v} := X_{\ell,i,v}$$
- For instruction $\ell : \text{if } X_i \text{ then goto } \ell' \text{ else } \ell''$,
 C_ℓ modifies the variables $X_{\ell',i,v}$ and $X_{\ell'',i,v}$ as follows:

$$X_{\ell',i,0} := X_1; \quad \triangleright 0 \text{ is definitely an error here}$$

$$X_{\ell',i,1} := X_{\ell,i,1}; \quad \triangleright 1 \text{ is an error if it was so before}$$

$$X_{\ell'',i,1} := X_1; \quad \triangleright 1 \text{ is definitely an error here}$$

$$X_{\ell'',i,0} := X_{\ell,i,0} \quad \triangleright 0 \text{ is an error if it was so before}$$
 For all $j \neq i$, we simply have, for both values of v ,

$$X_{\ell',i,v} := X_{\ell,i,v};$$

$$X_{\ell'',i,v} := X_{\ell,i,v}$$

Finally, these commands are put together to make the program p :

```
INIT;
loop X1 { choose C1 or C2 or ... Cm };
X2 := X0,1,0 * X0,2,0 * ... * X0,k,0
```

The outcome of the reduction is $(p, 2, 2k)$. Thus, we are asking whether the final value of X_2 depends on x with a degree of $2k$ or more, which will only happen if the loop is completed with x^2 in each of the variables $X_{0,i,0}$. This means that program b could be successfully simulated up to a point where the location was 0 and all variables *false* (that is, the program really terminated).

A comment is due on the subject of the loop bound (which we chose, somewhat arbitrarily, to be x). Obviously we do not know in advance how long b runs. But since it is input-free, if

b terminates, its running time is a constant and the desired output (x^{2k}) will be achieved for all values of x large enough to permit the complete simulation of b . If b does not terminate, the output will always be bounded by x^{2k-1} . \square

What is surprising in the above proof is the simplicity of the programs; curiously, they do not contain implicit data-flow in loops, which was the main challenge in the problem's solution.

11. CHECKING MULTIVARIATE BOUNDS

In this section we are moving to multivariate bounds. We show that checking a multivariate bound can be reduced to the univariate problem. We focus on the following decision problem (first defined in Section 3).

Problem 3.1. The *multivariate bound decision problem* is: Given a core-language program P and a monomial $m(x_1, \dots, x_n) = x_1^{d_1} \dots x_n^{d_n}$, report whether m constitutes an attainable bound on the final value of X_n ; namely whether constants \mathbf{x}_0 , $c > 0$ exist such that

$$\forall \mathbf{x} \geq \mathbf{x}_0 \exists \mathbf{x}' . \mathbf{x} \llbracket P \rrbracket \mathbf{x}' \wedge x'_n \geq cm(\mathbf{x}).$$

We later discuss how to solve the related search problem, asking to find a complete set of attainable bounds. If there is no polynomial upper bound on x'_n , then all monomials are attainable. This case can be detected in polynomial time using the algorithm of [BJK08], so we assume henceforth that x'_n is polynomially bounded. In fact, as already discussed, we may assume that all super-polynomially growing variables have been excluded.

Given P , we consider the set of attainable monomials (forming positive instances of the problem). We represent a monomial by a column vector of degrees: $\mathbf{d} = \langle d_1, \dots, d_n \rangle^T$. The main idea in this section is to make use of the *geometry* of this set of vectors by viewing the problem as a *linear programming* problem. Before proceeding, we recall some background knowledge.

11.1. Polyhedra. We recall some useful definitions and properties, all can be found in [Sch86].

Point $\mathbf{x} \in \mathbb{Q}^n$ is a *convex combination* of points $\mathbf{x}_1, \dots, \mathbf{x}_m$ if $\mathbf{x} = \sum_{i=1}^m a_i \cdot \mathbf{x}_i$, where all $a_i \geq 0$ and $\sum_i a_i = 1$. The *convex hull* of a set of points is the set of all their convex combinations.

A *rational convex polyhedron* $\mathcal{P} \subseteq \mathbb{Q}^n$ (*polyhedron* for short) can be defined in two equivalent ways:

- (1) As the set of solutions of a set of inequalities $A\mathbf{x} \leq \mathbf{b}$, namely $\mathcal{P} = \{\mathbf{x} \in \mathbb{Q}^n \mid A\mathbf{x} \leq \mathbf{b}\}$, where $A \in \mathbb{Q}^{m \times n}$ is a rational matrix of n columns and m rows, $\mathbf{x} \in \mathbb{Q}^n$ and $\mathbf{b} \in \mathbb{Q}^m$ are column vectors of n and m rational values respectively. Each linear inequality (specified by a row of A and the corresponding element of \mathbf{b}) is known as a *constraint*.
- (2) As the convex hull of a finite set of points \mathbf{x}_i and rays \mathbf{y}_j :

$$\mathcal{P} = \text{convhull}\{\mathbf{x}_1, \dots, \mathbf{x}_m\} + \text{cone}\{\mathbf{y}_1, \dots, \mathbf{y}_t\}, \quad (11.1)$$

or more explicitly: $\mathbf{x} \in \mathcal{P}$ if and only if $\mathbf{x} = \sum_{i=1}^m a_i \cdot \mathbf{x}_i + \sum_{j=1}^t b_j \cdot \mathbf{y}_j$ for some rationals $a_i, b_j \geq 0$, where $\sum_{i=1}^m a_i = 1$. The vectors $\mathbf{y}_1, \dots, \mathbf{y}_t$ are recession directions of \mathcal{P} , i.e., directions in which the polyhedron is unbounded; in terms of the constraint

representation, they satisfy $A\mathbf{y}_i \leq \mathbf{0}$. The points \mathbf{x}_i of a minimal set of generators are the *vertices* of \mathcal{P} .

For any set $\mathcal{S} \subseteq \mathbb{Q}^n$ we let $I(\mathcal{S})$ be $\mathcal{S} \cap \mathbb{Z}^n$, i.e., the set of integer points of \mathcal{S} . The *integer hull* of \mathcal{S} , commonly denoted by \mathcal{S}_I , is defined as the convex hull of $I(\mathcal{S})$. A polyhedron \mathcal{P} is *integral* if $\mathcal{P} = \mathcal{P}_I$. This is equivalent to stating that it is generated (Eq. 11.1) by a set of integer points and rays. In particular, its vertices are integral.

We will be using some results from [Sch86], regarding the *computational complexity* of algorithms on polyhedra. Thus we adhere to definitions used by the author, which we now recite for completeness. Schrijver denotes the bit-size of an integer x by $\|x\| = 1 + \lceil \log(|x| + 1) \rceil$; the bit-size of an n -dimensional vector \mathbf{a} is defined as $\|\mathbf{a}\| = n + \sum_{i=1}^n \|a_i\|$; and the bit-size of an inequality $\mathbf{a} \cdot \mathbf{x} \leq c$ as $1 + \|c\| + \|\mathbf{a}\|$. For a polyhedron $\mathcal{P} \subseteq \mathbb{Q}^n$ defined by $A\mathbf{x} \leq \mathbf{b}$, the *facet size*, denoted by $\|\mathcal{P}\|_\phi$, is the smallest number $\phi \geq n$ such that \mathcal{P} may be described by some $A\mathbf{x} \leq \mathbf{b}$ where each inequality in $A\mathbf{x} \leq \mathbf{b}$ fits in ϕ bits. The *vertex size*, denoted by $\|\mathcal{P}\|_\psi$, is the smallest number $\psi \geq n$ such that \mathcal{P} has a generator representation in which each of \mathbf{x}_i and \mathbf{y}_j fits in ψ bits (the size of a vector is calculated as above). The following theorem [Sch86, part of Theorem 10.2] relates the two measures:

Theorem 11.1. *Let \mathcal{P} be a rational polyhedron in \mathbb{Q}^n ; then $\|\mathcal{P}\|_\phi \leq 4n^2\|\mathcal{P}\|_\psi$.*

11.2. Monomial bounds and linear programming. From this point on, we fix \mathcal{S} to mean the set of attainable monomials of a given program (as vectors in \mathbb{N}^n).

Let $\mathbf{u} = \langle u_1, \dots, u_n \rangle \in \mathbb{N}^n$ be a (row) vector representing a univariate input, as in Section 7. Let $m(x_1, \dots, x_n) = x_1^{d_1} \dots x_n^{d_n}$. Instead of $x_1^{d_1} \dots x_n^{d_n}$ we write, concisely, $\mathbf{x}^{\mathbf{d}}$. Then note that

$$m(x^{u_1}, \dots, x^{u_n}) = x^{d_1 u_1 + \dots + d_n u_n} = x^{\mathbf{u} \cdot \mathbf{d}}.$$

We now state

Lemma 11.2. *Finding a tight upper bound on the final value of X_n given an initial univariate state \mathbf{u} is equivalent to the following optimization problem:*

$$\text{Given } \mathbf{u}, \text{ maximize } \mathbf{u} \cdot \mathbf{d} \quad \text{subject to } \mathbf{d} \in \mathcal{S}. \quad (11.2)$$

Note that the multivariate problem is simply to decide \mathcal{S} . The statement is by no means trivial. Luckily it comes easily out of the results of [BAH20].

Proof. According to [BAH20], given that the values computed by the program are polynomially bounded, the set \mathcal{S} of all the attainable *multivariate* monomials provides tight upper bounds as well (more precisely, taking the maximum of these monomials gives a well-defined function—since the set is finite—and this function is an asymptotic upper bound). If we plug \mathbf{u} as initial state, the set of bounds becomes $x^{\mathbf{u} \cdot \mathbf{d}}$ where \mathbf{d} ranges over \mathcal{S} . These are univariate monomials, so they are fully ordered. The highest bound, namely the maximum value of $\mathbf{u} \cdot \mathbf{d}$, is the tight worst-case upper bound for the program. \square

So far, we have only considered bounds which are polynomials, where the exponents are integer. However, a function $m(x_1, \dots, x_n) = x_1^{d_1} \dots x_n^{d_n}$ where the d_i are non-negative *rational numbers* is a perfectly valid candidate for comparison with the results of a computation. Let \mathcal{T} be the set of rational-valued vectors \mathbf{d} such that the corresponding function m is attainable. Clearly $\mathcal{S} \subseteq \mathcal{T}$; in fact, $\mathcal{S} = I(\mathcal{T})$. The inclusion is clearly strict. However, when we consider the optimization problem (11.2), we have




Figure 4: Assuming that the black, blue, red and green dots represent extremal elements of \mathcal{S} , the set \mathcal{T} is represented by the convex hull of the union of four boxes (shaded area in the left-hand drawing). The right-hand drawing illustrates \mathcal{T}' .

Lemma 11.3. For any vector $\mathbf{u} \in \mathbb{N}^n$,

$$\max_{\mathbf{d} \in \mathcal{T}} \mathbf{u} \cdot \mathbf{d} = \max_{\mathbf{d} \in \mathcal{S}} \mathbf{u} \cdot \mathbf{d}.$$

Proof. Since $\mathcal{S} \subseteq \mathcal{T}$, clearly $\max_{\mathbf{d} \in \mathcal{T}} \mathbf{u} \cdot \mathbf{d} \geq \max_{\mathbf{d} \in \mathcal{S}} \mathbf{u} \cdot \mathbf{d}$. Suppose that the inequality is strict; then there is an attainable function $m = \mathbf{x}^{\mathbf{d}}$ (with $\mathbf{d} \in \mathcal{T}$) such that $m(\mathbf{u})$ has a greater exponent than $m'(\mathbf{u})$, for all $m' = \mathbf{x}^{\mathbf{d}'}$ where $\mathbf{d}' \in \mathcal{S}$. This contradicts Corollary 4.2. \square

We now make a couple of observations on the shape of \mathcal{T} .

Lemma 11.4. For any $\mathbf{d} \in \mathcal{T}$, all integer points of the box $B_{\mathbf{0}, \mathbf{d}} = \{\mathbf{x} \mid \mathbf{0} \leq \mathbf{x} \leq \mathbf{d}\}$ are also in \mathcal{T} .

Proof. Immediate since if $\langle a_1, \dots, a_n \rangle \leq \langle b_1, \dots, b_n \rangle$ then $x^{a_1} \dots x^{a_n} \leq x^{b_1} \dots x^{b_n}$. \square

We conclude that \mathcal{T} is a union of boxes. We can say more.

Lemma 11.5. If functions $f, g : \mathbb{N}^n \rightarrow \mathbb{R}$ are attainable, so is $\max(f, g)$.

This can be easily checked against the definition (Definition 3.2).

Lemma 11.6. \mathcal{T} is convex.

Proof. Consider a monomial $m = \mathbf{x}^{\mathbf{d}}$ where \mathbf{d} is the convex combination $\sum c_j \cdot \mathbf{d}_j$ of some $\mathbf{d}_j \in \mathcal{T}$, where $\sum_{j=1}^n c_j = 1$. We refer to the i th component of \mathbf{d}_j as d_{ji} . Thus

$$m(\mathbf{x}) = \mathbf{x}^{\sum_j c_j \mathbf{d}_j} = \prod_i \prod_j x_i^{c_j d_{ji}} = \prod_j \prod_i x_i^{c_j d_{ji}} = \prod_j (\mathbf{x}^{\mathbf{d}_j})^{c_j}.$$

By the rational-weight form of the classic inequality of means [Ste04, Eq. (2.7)], a weighted geometric mean is bounded by the corresponding weighted arithmetic mean:

$$\prod_j (\mathbf{x}^{\mathbf{d}_j})^{c_j} \leq \sum_j c_j (\mathbf{x}^{\mathbf{d}_j}) \leq \max_j (\mathbf{x}^{\mathbf{d}_j});$$

The latter is an attainable function, therefore so is m . Hence $\mathbf{d} \in \mathcal{T}$. \square

So \mathcal{T} is convex, includes \mathcal{S} , and the maximum of $\mathbf{u} \cdot \mathbf{d}$ for any $\mathbf{u} > \mathbf{0}$ is obtained at a point of \mathcal{S} . We conclude that \mathcal{T} is an integral polyhedron [Sch86, §16.3]. It equals the convex hull of the boxes whose upper right corners are points of \mathcal{S} (Figure 11.2, left). It will be convenient in the next proof to extend \mathcal{T} by allowing negative numbers as well. This gives a polyhedron \mathcal{T}' which is unbounded in the negative direction (Figure 11.2, right); technically, $\mathcal{T}' = \mathcal{T} \cup \text{cone}(\langle -1, 0, \dots, 0 \rangle, \langle 0, -1, \dots, 0 \rangle, \dots, \langle 0, \dots, 0, -1 \rangle)$.

Now we can use known results on the complexity of linear programming to obtain our result.

Lemma 11.7. *The facet size of \mathcal{T}' is bounded by $4n^2(n + \|d_{max}\|)$, where d_{max} is the maximal attainable degree.*

This follows directly from Theorem 11.1, using the fact that our polyhedron is integral, i.e., its vertices are members of \mathcal{S} .

Theorem 11.8. *The multivariate-bound decision problem can be solved in space polynomial in the size of the given program and the maximal attainable degree d_{max} .*

Proof. As explained above, our approach is to show that we can decide whether $\mathbf{d} \in \mathcal{T}$. As PSPACE is closed under complement, we can consider the converse question: whether $\mathbf{d} \notin \mathcal{T}'$. This is equivalent to asking: is there a constraint $\mathbf{a} \cdot \mathbf{x} \leq b$, satisfied by \mathcal{T}' , such that $\mathbf{a} \cdot \mathbf{d} > b$. Our decision procedure is non-deterministic and *guesses* \mathbf{a} . Then we need to find b , which amounts to maximizing $\mathbf{a} \cdot \mathbf{x}$ over $\mathbf{x} \in \mathcal{T}'$. Note that \mathbf{a} will contain no negative numbers, since \mathcal{T}' does not satisfy any constraint with negative coefficients. In case \mathbf{a} has non-integer rational numbers, we can scale by their common denominator. Thus, w.l.o.g. we may assume \mathbf{a} to be integer-valued, and by Lemmas 11.2 and 11.3, finding whether b exists is equivalent to solving the univariate decision problem (Problem 3.5 on p. 8) with initial state \mathbf{a} , querying the degree $\mathbf{a} \cdot \mathbf{d}$.

So, we have reduced the multivariate-bound decision problem to the univariate problem. What is the complexity of the resulting algorithm? Assuming that we know d_{max} , we can impose the bound given by Lemma 11.7. Since \mathbf{a} consists of polynomially-big numbers (in terms of their bit size), our univariate algorithm (Section 7) solves the problem in polynomial space as stated.

What if we do not know d_{max} ? In this case we can search for it; we give the procedure below. □

In the case that d_{max} is not known, our complexity result is less satisfying when the monomial presented to the decision problem is of a degree much smaller than d_{max} . However, if what we really seek is not the decision problem for its own sake, but the discovery of tight upper bounds, then we do want to reach d_{max} , and we can then describe our algorithm's complexity as polynomial space in the *output-sensitive sense* (that is, depending on a quantity which is not the size of the input but of the output). First, let us state the complexity of the search procedure.

Theorem 11.9. *Given a core-language program P , the index of a chosen variable X_j , we can maximize $d_{max} = \sum_i d_i$ over all monomials $\mathbf{x}^{\mathbf{d}}$ that constitute an attainable lower bound on the final value of X_n . The algorithm's space complexity is polynomial in the size of P plus $\|d_{max}\|$.*

Proof. Recall that a polynomial-space non-deterministic algorithm can be determinized in polynomial space (Savitch's Theorem). We search for d_{max} using a deterministic implementation of our non-deterministic algorithm (Section 7). This deterministic algorithm can be

used to give a definite answer to the query: is a given degree d attainable for X_n . We start with $d = 1$ and repeatedly increase d (reusing the memory) until we hit a negative answer (as usual we assume that the possibility of super-polynomial growth has been eliminated first using [BJK08]). We can actually be quicker and only loop through powers of two, since this would give an approximation to d_{max} which is close enough to be used in Theorem 11.8. \square

12. GENERATING MULTIVARIATE BOUNDS DIRECTLY

Our analysis algorithm from Section 7 can actually be modified to yield multivariate bounds directly, as we show next. Given this statement, the reader may wonder why we bothered with the version for univariate bounds, and the reduction of multi-to-uni-variate. Let us reveal right away, then, that despite the simplicity of the following algorithm, we do not have a direct *proof* for its correctness and we approach it based on the previous algorithms.

The new algorithm is obtained thanks to a rather simple observation about the analyzer EV_M , namely, that the only operator applied to our (univariate) monomials is product. Since the product of multivariate monomials is again a (multivariate) monomial, we can replace the data type of our symbolic state \mathbf{s} from an n -tuple of *univariate* monomials to an n -tuple of *multivariate* ones. Such an n -tuple is a special case of AMP, and we also refer to it as *a monomial state* because it is used in the evaluator to abstract a state, just as the univariate state in the previous algorithm. Other than this change of datatype, no change to the code is necessary. In terms of implementation, a multivariate monomial is implemented as either an n -tuple of degrees, or a special constant representing 0. The space complexity of this representation remains polynomial in the number of variables and the maximum degree d_{max} . We call the evaluator, thus modified, EV_{MM} . To illustrate the idea, consider a 3-variable program. Then its initial state is represented by $Id = \langle x_1, x_2, x_3 \rangle$. Here are two examples of simple straight-line computations:

$$\begin{aligned} EV_{MM}[\mathbf{X}_3 := \mathbf{X}_2 * \mathbf{X}_1] Id &\Rightarrow (\langle x_1, x_2, x_1x_2 \rangle, E(3)) \\ EV_{MM}[\mathbf{X}_3 := \mathbf{X}_2 * \mathbf{X}_1 ; \mathbf{X}_2 := \mathbf{X}_3] Id &\Rightarrow^* (\langle x_1, x_1x_2, x_1x_2 \rangle, E(3)A(3, 2)) \end{aligned}$$

where the data-flow matrices $A(3, 2)$ and $E(3)$ are as in Section 7.2.

Proving the correctness of the algorithm falls, as usual, into soundness and completeness.

12.1. Soundness. The soundness for straight-line code is straight-forward to prove, so we content ourself with stating it:

Lemma 12.1. *Let \mathcal{C} be a loop-free command. Given any monomial state \mathbf{s} , if $EV_{MM}[\mathcal{C}]\mathbf{s} \Rightarrow^* (\mathbf{g}, M)$ (for any M) then $\exists \mathbf{p} \in \llbracket \mathcal{C} \rrbracket^S$ such that $\mathbf{g} \sqsubseteq \mathbf{p} \bullet \mathbf{s}$, and $M = \text{LDFM}(\mathbf{p})$.*

This means that if we set the initial state to Id , the monomials we generate are attainable in the sense of Corollary 4.2. We now wish to show the same property for loops. As in Section 9, the extension to EV_{MM}^* is immediate:

Lemma 12.2. *Let \mathcal{C} be a loop-free command. Given any monomial state \mathbf{s} , if $EV_{MM}^*[\mathcal{C}]\mathbf{s} \Rightarrow^* (\mathbf{g}, M)$ then $\exists \mathbf{p} \in Cl(\llbracket \mathcal{C} \rrbracket^S)$ such that $\mathbf{g} \sqsubseteq \mathbf{p} \bullet \mathbf{s}$, and $M = \text{LDFM}(\mathbf{p})$.*

As in Section 9, we proceed to cover any command by structural induction, where all loop-free commands are covered by the above lemmas, and similarly, given correctness for commands \mathcal{C} and \mathcal{D} , correctness for their sequential composition, or non-deterministic choice,

follow straight from definitions. So the main task is to prove the correctness for loops, more precisely the correctness of EV_{MM}^{\otimes} ; to this end we compare the possible results of $\text{GEN}(Z, \ell, \text{EV}_{MM}^*[\llbracket \mathbf{C} \rrbracket](\mathbf{s}[Z \leftarrow 0]))$ with our reference result $\text{Gen}(\text{Cl}(\llbracket \mathbf{C} \rrbracket^S))[x_\ell/\tau]$.

The soundness part of the proof repeats the arguments used with EV_M almost verbatim.

Lemma 12.3 (Soundness of GEN). *For any monomial state \mathbf{s} , and $Z \subset [n]$, if $(\mathbf{g}, M) \in \text{GEN}(Z, \ell, \text{EV}_{MM}^*[\llbracket \mathbf{C} \rrbracket](\mathbf{s}[Z \leftarrow 0]))$, then there is $\mathbf{r} \in \text{Gen}(\text{Cl}(\llbracket \mathbf{C} \rrbracket^S))$ such that $\mathbf{g} \sqsubseteq (\mathbf{r}[x_\ell/\tau]) \bullet \mathbf{s}$, and $\text{LDFM}(\mathbf{r})[x_\ell/\tau] = M$.*

Proof. Let $(\mathbf{s}', M') = \text{EV}_{MM}^*[\llbracket \mathbf{C} \rrbracket](\mathbf{s}[Z \leftarrow 0])$. When generalization is not applied, i.e., $\mathbf{g} = \mathbf{s}'$ and $M = M'$, we have, by inductive hypothesis, $\mathbf{p} \in \text{Cl}(\llbracket \mathbf{C} \rrbracket^S)$ such that $\mathbf{s}' = \mathbf{p} \bullet (\mathbf{s}[Z \leftarrow 0]) \sqsubseteq \mathbf{p} \bullet \mathbf{s}$ and $M' = \text{LDFM}(\mathbf{p})$, which is the same as $\text{LDFM}(\mathbf{r})[x_\ell/\tau]$ (since there is no τ).

Next, suppose that matrix M' satisfies the condition for generalization ($\forall i \notin Z : M'_{ii} = 1$), and let \mathbf{p} be an element of $\text{Cl}(\llbracket \mathbf{C} \rrbracket^S)$ such that $\mathbf{s}' \sqsubseteq \mathbf{p} \bullet \mathbf{s}[Z \leftarrow 0]$ and $M' = \text{LDFM}(\mathbf{p})$. Consider an index k to which generalization applies, i.e., $k \in Z$ and $M'_{kk} = 1$. Thus x_k is self-dependent in \mathbf{p} . Let $\mathbf{p}_Z = \mathbf{p} \bullet \text{Id}[Z \leftarrow 0]$. Then \mathbf{p}_Z is iterative; in fact, $\mathbf{p}_Z \sqsubseteq (\mathbf{p})$. Moreover, since $k \in Z$, $\mathbf{p}_Z[k] \sqsubseteq q$ such that $(\mathbf{p})[k] = x_k + q$. Now, in the Closure Algorithm, function Gen multiplies the polynomial q by τ (which is then substituted with x_ℓ), so multiplying $\mathbf{p}_Z[k] \bullet \mathbf{s}$ by $\mathbf{s}'[\ell]$, as GEN does, produces

$$\mathbf{s}^{\text{gen}}[k] \sqsubseteq (\mathbf{p}_Z[k] \bullet \mathbf{x}) \cdot \mathbf{s}'[\ell] \sqsubseteq ((\mathbf{p})^\tau[k][x_\ell/\tau]) \bullet \mathbf{x}.$$

We take \mathbf{r} to be $(\mathbf{p})^\tau$. Correctness of M is as in the previous proofs. \square

Once we can do the induction step for loops, we easily get

Lemma 12.4 (Soundness of EV_{MM}). *Let \mathbf{P} be a core-language program. If $\text{EV}_{MM}[\llbracket \mathbf{P} \rrbracket \text{Id}] \Rightarrow^* (\mathbf{g}, M)$ (for any M) then $\exists \mathbf{p} \in \llbracket \mathbf{P} \rrbracket^S$ such that $\mathbf{g} \sqsubseteq \mathbf{p}$.*

12.2. Completeness. We formulate our completeness claim based on the geometric view taken in the previous section. We maintain the convention that \mathcal{S} denotes the set of attainable monomials (asymptotic lower bounds on the final value of \mathbf{X}_n) for a program under discussion (arbitrary but fixed through the discussion), and $\mathcal{T}, \mathcal{T}'$ its extensions to rational-number polyhedra.

Now, in the previous subsection we have argued for soundness of EV_{MM} in the same way that we did for EV_M , that is by comparison to the Closure Algorithm, and established the relation $\mathbf{g} \sqsubseteq \mathbf{p} \bullet \mathbf{s}$ between EV_{MM} 's result and some $\mathbf{p} \in \llbracket \mathbf{C} \rrbracket^S$. A corresponding completeness claim might be that we can reach any monomial state $\mathbf{g} \sqsubseteq \mathbf{p} \bullet \mathbf{s}$. This is, however, not the case. Consider the two assignments:

$$\begin{aligned} \mathbf{X}_2 &:= \mathbf{X}_1 + \mathbf{X}_2; \\ \mathbf{X}_2 &:= \mathbf{X}_2 * \mathbf{X}_2 \end{aligned}$$

In EV_{MM} (starting with Id), after the first assignment we have two monomial states, $\langle x_1, x_1 \rangle$ and $\langle x_1, x_2 \rangle$; and after the second, we have $\langle x_1, x_1^2 \rangle$ and $\langle x_1, x_2^2 \rangle$. But the Closure Algorithm generates $\langle x_1, x_1^2 + x_1x_2 + x_2^2 \rangle$. Note that there is a monomial, x_1x_2 , which is not obtained by EV_{MM} . But this is OK, because the degree vector $(1, 1)$ lies on the line between $(2, 0)$ and $(0, 2)$; i.e., it is not a vertex of the attainable-monomial polyhedron. In order to guarantee that we find the vertices, it suffices to compare the results on univariate states—this is the main insight of Section 11.

Lemma 12.5. *Let \mathcal{C} be any command. Given any monomial state \mathbf{s} , any $\mathbf{p} \in \llbracket \mathcal{C} \rrbracket^S$, and any univariate state \mathbf{x} , there are a monomial state \mathbf{g} and matrix M such that*

$$\begin{aligned} \text{EV}_{MM} \llbracket \mathcal{C} \rrbracket \mathbf{s} &\Rightarrow^* (\mathbf{g}, M), \\ \mathbf{g} \bullet \mathbf{x} &\geq \mathbf{p} \bullet \mathbf{s} \bullet \mathbf{x} \text{ and} \\ M &= \text{LDFM}(\mathbf{p}). \end{aligned}$$

Proof. Fix \mathbf{s} , \mathbf{p} and \mathbf{x} . By Theorem 9.1 there is $\mathbf{y} \geq \mathbf{p} \bullet \mathbf{s} \bullet \mathbf{x}$ such that

$$\text{EV}_M \llbracket \mathcal{C} \rrbracket (\mathbf{s} \bullet \mathbf{x}) \Rightarrow^* (\mathbf{y}, \text{LDFM}(\mathbf{p})).$$

Now, since EV_{MM} simulates the execution of EV_M step by step, with the sole difference that it operates on the elements of \mathbf{s} instead of those of $\mathbf{s} \bullet \mathbf{x}$, and that the only operations applied to monomials are copying and product, it should be clear that

$$(\text{EV}_{MM} \llbracket \mathcal{C} \rrbracket \mathbf{s}) \bullet \mathbf{x}$$

evaluates to the same set of values as

$$\text{EV}_M \llbracket \mathcal{C} \rrbracket (\mathbf{s} \bullet \mathbf{x})$$

and this justifies the lemma. □

Corollary 12.6 (Completeness of EV_{MM}). *Let \mathcal{C} be a core-language program. Let $m = x_1^{d_1} \dots x_n^{d_n}$ be a monomial representing an attainable bound on the final value of x_n after executing \mathcal{C} , and such that its degree vector \mathbf{d} is a vertex of \mathcal{T}' . Then $\text{EV}_{MM} \llbracket \mathcal{C} \rrbracket Id \Rightarrow^* (\mathbf{g}, M)$ (for some M) such that $\mathbf{g}[n] = m$.*

Proof. Since \mathbf{d} is a vertex, there is a univariate state \mathbf{x} for which \mathbf{d} is the unique solution to the optimization problem (11.2), meaning that $m \bullet \mathbf{x}$ maximizes the n th component in $\mathbf{p} \bullet \mathbf{x}$ when \mathbf{p} ranges over $\llbracket \mathcal{C} \rrbracket^S$. By the last lemma, this can be matched by EV_{MM} , and since EV_{MM} is also sound (i.e., it only produces attainable monomials), it must produce m . □

Since we can find any vertex of the polyhedron, we have a complete solution in the sense that the **max** of all these monomials is a tight multivariate upper bound on the queried variable, as in Corollary 4.2.

12.3. Discussion. We have given a polynomial-space solution to the problem of *generating* a bound. Since the algorithm can generate all the vertices of the attainable-monomial polyhedron, it could be used to implement an alternative *decision* procedure as well, replacing the proof of Theorem 11.8. Note that, ultimately, both proofs rely on the insight expressed by Lemma 11.2, and even if this is a redundancy, we thought that presenting both approaches makes the picture more complete.

Looking back at EV_{MM} , as we have already noticed that the initial state does not really influence any decision by the algorithm, it is easy to see that we could rewrite the algorithm to remove this parameter. The algorithm then abstracts a command to a set in $\text{AMPo1} \times M_n(\mathbb{B})$, rather than a function that maps initial states to such sets. Next, we give this version in equational form, in the style of the Closure Algorithm. The result set may be of exponential size, so the non-deterministic interpreter form of the algorithm is useful to prove the PSPACE complexity class, but we feel that the following presentation is of a certain elegance. The abstraction of command \mathcal{C} is denoted $\llbracket \mathcal{C} \rrbracket^M$ (M for Monomials).

We define $\llbracket \mathbf{C} \rrbracket^M$, first for non-looping commands:

$$\begin{aligned}
\llbracket \text{skip} \rrbracket^M &= \{(Id, I)\} \\
\llbracket X_i := X_j \rrbracket^M &= \{(\text{set}_{ij}, A(j, i))\} \\
\llbracket X_i := X_j + X_k \rrbracket^M &= \{(\text{set}_{ij}, A(j, i) + A(k, i)), (\text{set}_{ik}, A(j, i) + A(k, i))\} \\
\llbracket X_i := X_j * X_k \rrbracket^M &= \{(\text{mul}_{ijk}, E(i))\} \\
\llbracket \text{choose } C_1 \text{ or } C_2 \rrbracket^S &= \llbracket C_1 \rrbracket^S \cup \llbracket C_2 \rrbracket^S \\
\llbracket C_1; C_2 \rrbracket^S &= \llbracket C_2 \rrbracket^S \odot \llbracket C_1 \rrbracket^S,
\end{aligned}$$

where $(\mathbf{p}, A) \odot (\mathbf{q}, B) \stackrel{def}{=} (\mathbf{p} \bullet \mathbf{q}, A \cdot B)$.

Finally,

$$\llbracket \text{loop } X_\ell \{C\} \rrbracket^M = Cl(Gen(\ell, Cl(\llbracket C \rrbracket^M)))$$

where $Cl()$ denotes closure under \odot , and

$$Gen(\ell, \mathcal{S}) = \bigcup_{\substack{(\mathbf{p}, M) \in \mathcal{S} \\ Z \subset [n]}} GEN(Z, \ell, (\mathbf{p} \bullet (Id[Z \leftarrow 0]), M)),$$

with function GEN defined just as in EV_{MM} (namely just as in EV_M , but calculating with multivariate monomials), as illustrated in the following example.

Example 12.7. Consider the following loop:

```

loop X5 {
  choose
    { X3 := X1; X4 := X2 }
  or
    X1 := X3 + X4
}

```

The abstraction of the *loop body* clearly yields the following pairs:

$$\begin{aligned}
&(\langle x_1, x_2, x_1, x_2, x_5 \rangle, \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}), \\
&(\langle x_3, x_2, x_3, x_4, x_5 \rangle, \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}), \\
&(\langle x_4, x_2, x_3, x_4, x_5 \rangle, \begin{bmatrix} 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix})
\end{aligned}$$

In the closure of the loop body's abstraction we will find compositions of these results; e.g., composing the first with the second (respectively the third) corresponds to simulating the instructions:

$$\begin{aligned} X_3 &:= X_1; X_4 := X_2; \\ X_1 &:= X_3 + X_4; \end{aligned}$$

and yields the results:

$$\begin{aligned} &(\langle x_1, x_2, x_1, x_2, x_5 \rangle, \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}), \\ &(\langle x_2, x_2, x_1, x_2, x_5 \rangle, \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 1 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}). \end{aligned}$$

To illustrate the effect of GEN , let $Z = \{1, 3, 4\}$. Inspecting the above matrix, we note that 3, 4 must be in the set to permit generalization, since they denote non-self-dependent entries; the presence of 1 in Z implies that we treat X_1 as an accumulator. The above two matrices produce, respectively, the following results

$$\begin{aligned} &(\langle 0, x_2, 0, x_2, x_5 \rangle, \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}), \\ &(\langle x_2x_5, x_2, 0, x_2, x_5 \rangle, \begin{bmatrix} 1 & 0 & 1 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}). \end{aligned}$$

We state the correctness of $\llbracket \cdot \rrbracket^M$, which is based on the results established for EV_{MM} , by the simple observation that

$$\mathbf{p} \in \llbracket \mathbf{C} \rrbracket^M \iff (EV_{MM} \llbracket \mathbf{C} \rrbracket Id \Rightarrow^* (\mathbf{p}, LDFM(\mathbf{p}))).$$

Theorem 12.8. *Let \mathbf{P} be a core-language program. The set $\llbracket \mathbf{P} \rrbracket^M$ satisfies the following properties:*

- (1) (soundness) *If $\mathbf{m} \in \llbracket \mathbf{P} \rrbracket^M$ then $\exists \mathbf{p} \in \llbracket \mathbf{P} \rrbracket^S$ such that $\mathbf{m} \sqsubseteq \mathbf{p}$.*
- (2) (completeness) *Let $m = x_1^{d_1} \cdots x_n^{d_n}$ be a monomial representing an attainable bound on the final value of variable X_n after executing \mathbf{P} , and such that its degree vector \mathbf{d} is a vertex of \mathcal{T}' . Then there is $\mathbf{m} \in \llbracket \mathbf{P} \rrbracket^M$ such that $\mathbf{m}[n] = m$.*

Note that the soundness clause indicates that \mathbf{m} is an attainable AMP for \mathbf{P} ; and that in the completeness clause, the choice of X_n is arbitrary. The same property holds for any chosen variable (and the corresponding set of attainable monomial bounds).

13. A FEW SIMPLE EXTENSIONS

A natural follow-up to the above results is to extend them to richer programming languages. In this section we briefly discuss four such extensions. First, two trivial ones:

- We can allow any polynomial (with positive coefficients) as the right-hand side of an assignment—this is just syntactic sugar.
- We can include the weak (non-deterministic) assignment form which says “let X be a natural number *bounded* by the following expression.” This construct may be useful if we want to use our language to approximately simulate the behavior of programs that have arithmetic expressions which we do not support, as long as a preprocessor can statically determine polynomial bounds on such expressions. No change to the algorithm is required.

Next, two extensions which are not trivial, but very easy.

Resets. [BA10] shows that it is possible to enrich the core language by a reset instruction, $X := 0$, and still solve the problem of distinguishing polynomially-bounded variables from potentially super-polynomial ones. Conceptually this is perhaps not a big change, but technically it caused a jump in the complexity of the solution, from PTIME to PSPACE-complete. Our tight-bound problem is already PSPACE-complete, and we can extend our solution to handle resets without a further increase in complexity. In fact, with the abstract-interpreter algorithm, adding the resets is very smooth. We add a special constant $\underline{0}$, which means *definitely zero* — the symbolic value of a variable that has been reset. This value is treated using the natural laws, i.e., $\underline{0} + x = x$, $\underline{0} * x = \underline{0}$. Resetting a variable cuts data-flow, so the abstraction of $X := 0$ reflects this in the LDFM:

$$EV_M \llbracket X_i := 0 \rrbracket \mathbf{s} \Rightarrow (\mathbf{s}[i \leftarrow \underline{0}], E(i))$$

Additional changes are necessary in the analysis of loops. If the loop bound is definitely zero, we know for certain that the loop body is not executed (unlike the general case, where skipping the loop is just a *possibility*). Thus we rewrite the corresponding section of our interpreter as follows:

$$EV_M \llbracket \text{loop } X_\ell \{C\} \rrbracket \mathbf{s} = \begin{cases} (\mathbf{s}, I) & \text{if } \mathbf{s}[\ell] = \underline{0}, \\ EV_M^{\otimes} \llbracket C \rrbracket (\ell, \mathbf{s}) & \text{otherwise.} \end{cases}$$

and similarly for EV_{MM} . Inside EV_M^{\otimes} (respectively EV_{MM}^{\otimes}), we have to deal with the case that a definitely-zero variable is selected for the set Z of variables to be set to zero. Let us recall that this set serves two purposes:

- (1) To mask out non-self-dependent variables. If the variable is initially definitely zero, there is of course no need to mask it out, and moreover, replacing $\underline{0}$ with an ordinary zero may cause imprecision, in case that an internal loop depends on that variable.
- (2) To isolate the increment q in computations that set $x'_i = x_i + q$, by putting $\underline{0}$ in x_i . This has to work, even in the case that x_i is initially $\underline{0}$; while, as above, replacing $\underline{0}$ with an ordinary zero would be wrong.

From these considerations, the correct code for Ev_M^{\otimes} becomes the following.

$$\begin{aligned} \text{Ev}_M^{\otimes}[\mathbf{C}](\ell, \mathbf{s}) &\Rightarrow (\mathbf{s}, I) \\ \text{Ev}_M^{\otimes}[\mathbf{C}](\ell, \mathbf{s}) &\Rightarrow \text{let } Z \subseteq \{i \mid \mathbf{s}[i] \neq \mathbf{0}\}; \\ &\quad \text{let } (\mathbf{s}_1, M_1) = \text{GEN}(Z, \ell, \text{Ev}_M^*[\mathbf{C}](\mathbf{s}[Z \leftarrow \mathbf{0}])); \\ &\quad \text{let } (\mathbf{s}_2, M_2) = \text{Ev}_M^{\otimes}[\mathbf{C}](\ell, \mathbf{s}_1) \\ &\quad \text{in } (\mathbf{s}_2, M_1 \cdot M_2) \end{aligned}$$

while in function GEN :

$$\begin{aligned} \text{GEN}(Z, \ell, (\mathbf{s}, M)) &= \text{if } (\forall i \notin Z : (\mathbf{s}[i] = \mathbf{0} \vee M_{ii} = 1)) \\ &\quad \text{then } (\mathbf{s}^{gen}, M^{gen}) \text{ else } (\mathbf{s}, M) \end{aligned}$$

where

$$\begin{aligned} \mathbf{s}^{gen} &= \mathbf{s} \prod_{i: M_{ii}=1 \wedge (i \in Z \vee \mathbf{s}[i]=\mathbf{0})} \mathbf{mul}_{iil} \\ M^{gen} &= (I \wedge M)M(I \wedge \overline{M}) \vee (I \wedge M) \end{aligned}$$

Unbounded “value”. In [BAP16] we proposed to extend our language with an instruction $\mathbf{X} := *$, sometimes called “havoc \mathbf{X} .” The intention is to indicate that \mathbf{X} is set to a value which is not expressible in the core language. This may be useful if we want to use our core-language programs as an abstraction of real programs, that possibly perform computations that we cannot express or bound by a polynomial. It may also be used to model a **while**-loop (as done, for example, in [JK09]): we do not attempt to analyze the loop for termination or even for bounds (this is outside the scope of our algorithm; such loops depend on concrete conditionals which we do not model at all). So all we suggest is to analyze what happens in the loop under the assumption that we cannot bound the number of iterations. This can be simulated by setting the loop bound to $*$.

The implementation is very similar to that of `reset`, and moreover, they can be combined. The rules for calculating with $*$ are: $* \cdot 0 = 0$, and $* \cdot m = *$ for any other monomial m . And, of course, if an output variable comes out as $*$ it is reported to have no polynomial upper bound.

14. RELATED WORK

Bound analysis, in the sense of finding symbolic bounds for data values, iteration bounds and related quantities, is a classic field of program analysis [Weg75, Ros89, LM88]. It is also an area of active research, with tools being currently (or recently) developed including COSTA [AAG⁺12], APROVE [GAB⁺17], CIAOPP [LDK⁺18], C^4B [CHS15], LOOPUS [SZV17] and CoFLoCo [FH14]—and this is just for imperative programs. There is also work on functional and logic programs, term rewriting systems, recurrence relations, etc. which we cannot attempt to survey here.

Our programming language, which is based on bounded loops, is similar to Meyer and Ritchie’s language of *Loop Programs* [MR67], which differed by being deterministic and including constants; it is capable of computing all the primitive recursive functions and too strong to obtain decidability results. Similar languages have been used as objects for analysis of growth-rate (in particular, identifying polynomial complexity) in [KA80, KN04, NW06,

JK09]. These works, too, considered deterministic languages for which precise analysis is impossible, but on reading them, one can clearly see that there are some clear limits to the aspects of the language that they analyze (such as using a loop counter as a bound, not relying on the assumption that the iteration count is always completed). Such considerations led to the definition of the weak, non-deterministic language in [BJK08]. Recent work in static analysis of fully-capable programs [BEF⁺16] combines a subsystem for loop-bound analysis (via ranking functions) with a subsystem for growth-rate analysis, which establishes symbolic bounds on data that grow along a loop, based on loop bounds provided by the loop analysis. This may be seen, loosely speaking, as a process that reduces general programs to bounded-loop programs which are then analyzed. Our previous paper [BAH20] was, however, the first that gave a solution for computing tight polynomial bounds which is complete for the class of programs we consider.

The problem of deciding termination of a weak programming language has been considered in several works; one well-studied example considers single-path while loops with integer variables, affine updates, and affine guard conditions [Tiw04, Bra06, HOW19]. Its decidability was first proved in [HOW19], but it is not shown to what complexity class this problem belongs. It would be interesting to see whether this can also be shown to be PSPACE-complete.

Dealing with a somewhat different problem, [MOS04, HOP⁺18] both check, or find, *invariants* in the form of polynomial equations, so they can identify cases where a final result is *precisely* polynomial in the input values. We find it remarkable that they give complete solutions for weak languages, where the weakness lies in the non-deterministic control-flow, as in our language. However, their language has no mechanism to represent bounded loops, whereas in our work, it is all about these loops.

15. CONCLUSION

We have complemented previous results on the computability of polynomial bounds in the weak imperative language defined in [BJK08]. While [BJK08] established that the existence of polynomial bounds is decidable, [BAH20] showed that tight bounds are computable. Here we have shown that they can be computed in polynomial space, and that this is optimal in the sense of PSPACE-hardness. We have thus settled the question of the complexity class of this program-analysis problem. Interestingly, this improvement required some new ideas on top of [BAH20], including the connection of multivariate bounds to univariate bounds based on the geometry of the set of monomial bounds, and the computation of bounds using an abstraction which is space-economical, including only degree vectors and data-flow matrices.

Some challenging open problems regarding bound computation remain:

- (1) Whether the problem is solvable at all for various extensions to the language (see [BAK12, BAH20] for further discussion).
- (2) How to compute more precise polynomial bounds (with explicit constant factors); can we make them tight?
- (3) Computation of tight super-polynomial bounds (currently, if a variable cannot be bounded by a polynomial, we give no upper bound).

Another problem we are interested in is: if one of our programs computes a polynomially-bounded result, how high can the degree of the bounding polynomial be, in terms of the size of the program? We conjecture that the degree is at most exponential in the program

size. However, a proof which we attempted (in an earlier technical report) turned out to be wrong, so the problem stands open.

Acknowledgement. We thank the anonymous referees, whose diligent reviews have been invaluable to our work on this article. On the occasion of his 80th birthday, we would also like to acknowledge the enormous contributions made by Neil D. Jones in this area, and his influence on our work.

APPENDIX A. JUSTIFICATION OF THE CLOSURE ALGORITHM

The purpose of this section is to clarify how the Closure Algorithm (Section 6) relates to the algorithm we published in [BAH20], and justify the correctness of the new version. The justification amounts to a proof that the algorithm computes a set of lower bounds which is equivalent to the set computed by [BAH20]; we are thus relying on the correctness proof of [BAH20].

A.1. Background. We begin with a stock of definitions and auxiliary results from [BAH20]. This subsection contains no novel material and is included for completeness.

Definition A.1 (Subsumption order). For τ -multivariate polynomials p, q we define $p \triangleleft q$ to hold if for every monomial \mathbf{m} of p some multiple of \mathbf{m} appears in q . We then say that p is subsumed by q .

For example, we have $x_1x_2 \triangleleft \tau x_1x_2 + x_3$. Moreover, $x_1 + \tau x_1 \triangleleft \tau x_1$. It should be clear that this order agrees with the asymptotic growth-rate order commonly denoted by $p = O(q)$. The definition is extended to MPs component-wise. Note that \triangleleft is a coarser relation than the fragment relation \sqsubseteq , i.e., $\mathbf{p} \sqsubseteq \mathbf{q} \Rightarrow \mathbf{p} \triangleleft \mathbf{q}$. For example, x_1x_2 is not a fragment of τx_1x_2 , but is subsumed by it.

Definition A.2. $\mathbf{p} \in \tau\text{MPol}$ is called *idempotent* if $\mathbf{p} \bullet \mathbf{p} = \mathbf{p}$.

Note that this is composition in the abstract domain. So, for instance, $\langle x_1, x_2 \rangle$ is idempotent, and so is $\langle x_1 + x_2, x_2 \rangle$, while $\langle x_1x_2, x_2 \rangle$ and $\langle x_1 + x_2, x_1 \rangle$ are not.

Let $\mathbf{p} \in \tau\text{MPol}$. Recall (Section 9) that its *dependence graph* $G(\mathbf{p})$ is a directed graph with node set $[n]$, which includes an arc $i \rightarrow j$ if and only if $\mathbf{p}[j]$ depends on x_i .

The following lemma (similar to Lemma 9.10) is from [BAH20].

Lemma A.3. *Suppose that the SDL \mathcal{S} is polynomially bounded, $\mathbf{p} \in \tau\text{MPol}$ is attainable over \mathcal{S} , and $\alpha(\mathbf{p})$ is idempotent. Then $G(\mathbf{p})$ does not have any simple cycle longer than one arc. In other words, $G(\mathbf{p})$ is quasi-acyclic.*

We assume that the loop under analysis is polynomially bounded; therefore we can rely on the above property of $G(\mathbf{p})$, being quasi-acyclic. We assume, w.l.o.g., that the variables are indexed in an order consistent with $G(\mathbf{p})$, so that if x_i depends on x_j then $j \leq i$. We refer to an (abstract) MP satisfying this as *neat*.

A.2. The Old Closure Algorithm. We now recall the algorithm from [BAH20], to facilitate its comparison with the current version. We are only considering the SDL analysis component—the rest is identical. The following function from [BAH20] is equivalent in role to function LC (Equation 6.2):

SOLVE(\mathcal{S})

- (1) Set $T = \mathcal{S}$.
- (2) Repeat the following steps until T remains fixed:
 - (a) Closure: Set T to $Cl(T)$.
 - (b) Generalization: For all $\mathbf{p} \in T$ such that $\mathbf{p} \bullet \mathbf{p} = \mathbf{p}$, add \mathbf{p}^τ to T .

The reader may have noticed that in this algorithm, generalization is applied to idempotent elements, instead of iterative elements. As an idempotent element is not necessarily iterative, the definition of generalization we used (Definition 6.15) does not suffice.

Definition A.4 (extended generalization). Let $\mathbf{p} \in \tau\text{AMPol}$; define \mathbf{p}^τ by:

$$\mathbf{p}^\tau[i] = \begin{cases} \mathbf{p}[i] + \tau\mathbf{p}[i]'' & \text{if } i \in \text{SD}(\mathbf{p}) \\ \mathbf{p}[i] & \text{otherwise,} \end{cases}$$

where $\mathbf{p}[i]''$ is defined as the sum of all iterative monomials of $\mathbf{p}[i]$, except for x_i (a monomial in \mathbf{p} is iterative if it consists of variables self-dependent in \mathbf{p}). We employ the unmotivated notation $\mathbf{p}[i]''$ just for consistency with [BAH20]).

Note that the addition in this definition is abstract (see example below). It is easy to verify that Definition 6.15 agrees with this definition (up to subsumption) if \mathbf{p} is iterative (and, in particular, τ -free; monomials that include τ are not considered iterative).

Example A.5. We illustrate Definition A.4 with an example. Let $\mathbf{p} = \langle x_1, x_2, x_2, x_4 + \tau x_2 + \tau x_1 + x_1 x_2 \rangle$. Then $\text{SD}(\mathbf{p}) = \{1, 2, 4\}$. This τ -AMP is not iterative, because of the τ -monomials in $\mathbf{p}[4]$. However, it is idempotent, and applying Definition A.4, $\mathbf{p}^\tau = \langle x_1, x_2, x_2, x_4 + \tau x_2 + \tau x_1 + \tau x_1 x_2 \rangle$. Note also that \mathbf{p} is neat: $G(\mathbf{p})$ is shown in Figure 5.




Figure 5: Dependence graph for the example MP. Nodes are labeled x_j rather than just j for readability.

Example A.6. Next, we give an example for a complete computation of SOLVE . We consider the loop:

```
loop X3 { X1 := X1 + X2; X2 := X2 + X3; X4 := X3 }
```

The body of the loop is evaluated symbolically and yields the abstract multi-polynomial:

$$\mathbf{p} = \langle x_1 + x_2, x_2 + x_3, x_3, x_3 \rangle$$

This singleton is the initial set T . To obtain $Cl(T)$, we include the identity MP, and also compute the compositions

$$\begin{aligned} (\mathbf{p})^{\bullet(2)} &= \mathbf{p} \bullet \mathbf{p} = \langle x_1 + x_2 + x_3, x_2 + x_3, x_3, x_3 \rangle; \\ (\mathbf{p})^{\bullet(3)} &= (\mathbf{p})^{\bullet(2)}. \end{aligned}$$

Here the closure computation stops. Next, in the generalization phase, since $\mathbf{p}^{\bullet(2)}$ is idempotent, we compute:

$$\mathbf{q} = ((\mathbf{p})^{\bullet(2)})^\tau = \langle x_1 + \tau x_2 + \tau x_3, x_2 + \tau x_3, x_3, x_3 \rangle$$

So, the following set results:

$$\begin{aligned} \langle x_1, x_2, x_3, x_4 \rangle, & \qquad \qquad \qquad \langle x_1 + x_2, x_2 + x_3, x_3, x_3 \rangle \\ \langle x_1 + x_2 + x_3, x_2 + x_3, x_3, x_3 \rangle, & \qquad \langle x_1 + \tau x_2 + \tau x_3, x_2 + \tau x_3, x_3, x_3 \rangle \end{aligned}$$

applying closure again, we obtain some additional results:

$$\begin{aligned} \mathbf{q} \bullet \mathbf{p} &= \langle x_1 + x_2 + x_3 + \tau x_2 + \tau x_3, x_2 + x_3 + \tau x_3, x_3, x_3 \rangle \\ (\mathbf{q})^{\bullet(2)} &= \langle x_1 + \tau x_2 + \tau x_3 + \tau^2 x_3, x_2 + \tau x_3, x_3, x_3 \rangle \\ (\mathbf{q})^{\bullet(2)} \bullet \mathbf{p} &= \langle x_1 + x_2 + x_3 + \tau x_2 + \tau x_3 + \tau^2 x_3, x_2 + x_3 + \tau x_3, x_3, x_3 \rangle \end{aligned}$$

The last element is idempotent but applying generalization does not generate anything new. Thus the final set T has been reached. The reader may reconsider the source code to verify that we have indeed obtained tight bounds for the loop. Note that \mathbf{p} itself is not idempotent, but it is iterative. Our improved version of the Closure Algorithm (Section 6) would generalize \mathbf{p} , adding $\mathbf{p}^\tau = \langle x_1 + \tau x_2, x_2 + \tau x_3, x_3, x_3 \rangle$. This result is subsumed by \mathbf{q} above, showing that the final sets of bounds resulting from the two algorithms are equivalent.

Our new algorithm is simpler than the old one as it performs only one round of generalization—it does not “iterate until fixed-point” like SOLVE. This is a simplification, since we do not have to consider the issue of correctly generalizing MPs that already include τ . It can be proved that every idempotent AMP has an iterative kernel; but many AMPs which are not idempotent also have an iterative kernel. One could say that the new algorithm is more eager to generalize. The motivation for this change was to eliminate the idempotence condition: in order to judge that an AMP is idempotent we need full information about it. Our PSPACE algorithm operates with very partial information and we derive its soundness by establishing that what it generalizes is an iterative fragment of *some* AMP in the closure, which we cannot guarantee to be idempotent. Thus, we made this change in order to justify the PSPACE algorithm. Nonetheless, we believe that our current variant may be more efficient in practice than the old one, mostly because of the elimination of the outer “iterate until fixed point” loop. It would be our choice, should we prefer to implement a closure-based algorithm (perhaps because determinization of the PSPACE algorithm, which is non-deterministic, would result in a deterministic algorithm worse in practice).

A.3. Equivalence of the algorithms. We now prove that our revised algorithm gives results equivalent to the former version. Basically, we prove subsumption of the two sets of results in each other.

Lemma A.7. *Every $\mathbf{p} \in \text{SOLVE}(\mathcal{S})$ is subsumed by a member of $LC(\mathcal{S})$.*

Proof. We prove inductively, based on the structure of the algorithm, that the set T , computed by SOLVE, is subsumed by $LC(\mathcal{S})$.

Base case: $T = \mathcal{S}$ is, of course, contained in \mathcal{S} .

Induction Step 1: Setting T to $Cl(T)$ preserves the relation because $LC(\mathcal{S})$ is closed under composition, as one can clearly see from its definition.

Induction Step 2: This step adds \mathbf{p}^τ to T for idempotent AMPs \mathbf{p} . We shall now prove that $\mathbf{p}^\tau \triangleleft LC(\mathcal{S})$. We consider two sub-cases. If \mathbf{p} is τ -free, we infer that $\mathbf{p} \in Cl(\mathcal{S})$. Then it is clearly also in $LC(\mathcal{S})$. If \mathbf{p} has τ 's, we observe that $\mathbf{p}[1/\tau]$ is what one gets by performing the same *composition* operations that yielded \mathbf{p} , but skipping the generalizations. So $\mathbf{p}[1/\tau] \in Cl(\mathcal{S}) \subseteq LC(\mathcal{S})$. It is easy to see that $\mathbf{p}[1/\tau]$ is idempotent as well; thus $(\mathbf{p}[1/\tau])^\tau$ is formed by function *Gen*. Also, by IH, there is $\mathbf{q} \in LC(\mathcal{S})$ such that $\mathbf{p} \triangleleft \mathbf{q}$. Using a few lemmas proved below, we conclude that $\mathbf{q} \bullet (\mathbf{p}[1/\tau])^\tau \bullet \mathbf{p}[1/\tau] \triangleright \mathbf{p}^\tau$; specifically,

$$\begin{aligned} \mathbf{q} \bullet (\mathbf{p}[1/\tau])^\tau \bullet \mathbf{p}[1/\tau] &\triangleright \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau \bullet \mathbf{p}[1/\tau] && \text{(Lemma A.8)} \\ &\triangleright \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau && \text{(Lemma A.11)} \\ &\triangleright \mathbf{p}^\tau && \text{(Lemma A.9)} \end{aligned}$$

The above product is in $LC(\mathcal{S})$, as we have established that each of its factors is present in the set $Gen(Cl(\mathcal{S}))$; we conclude $\mathbf{p}^\tau \triangleleft LC(\mathcal{S})$. \square

Lemma A.8. *For $p, q \in \tau\text{APo1}$ and $\mathbf{r} \in \tau\text{AMPo1}$, if $\mathbf{q} \triangleright \mathbf{p}$ then $\mathbf{q} \bullet \mathbf{r} \triangleright \mathbf{p} \bullet \mathbf{r}$.*

Proof. Let \mathbf{m} be a monomial of $\mathbf{p} \bullet \mathbf{r}$. It is obtained by substituting components $\mathbf{r}[i]$ into variables x_i of a monomial \mathbf{p} of p . Now, q is known to contain some multiple \mathbf{q} of \mathbf{p} , and clearly substituting \mathbf{r} in \mathbf{q} yields a multiple of \mathbf{m} . \square

Lemma A.9. *Let $\mathbf{p} \in \tau\text{AMPo1}$ be idempotent and neat. Then $\mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau \triangleright \mathbf{p}^\tau$.*

Proof. We prove, by induction on i , that the inequality holds for the first i components of the respective AMPs; we write this using “array slice notation” as $(\mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau)[1..i] \triangleright \mathbf{p}^\tau[1..i]$.

Base case: Since \mathbf{p} is neat we know that $\mathbf{p}[1] = x_1$. It is also clear that $(\mathbf{p}[1/\tau])^\tau[1] = x_1$. Consequently, $(\mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau)[1] = x_1 = \mathbf{p}^\tau[1]$.

Induction step: let $i > 1$. If $i \notin \text{SD}(\mathbf{p})$, then $\mathbf{p}[i]$ only depends on variables x_j with $j < i$. This allows us to use the induction hypothesis:

$$\begin{aligned} \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau[i] &= \mathbf{p} \bullet \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau[i] && \text{by idempotence} \\ &= \mathbf{p}[i] \bullet \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau && \\ &\triangleright \mathbf{p}[i] \bullet \mathbf{p}^\tau && \text{by IH} \\ &\triangleright \mathbf{p}[i] \bullet \mathbf{p} = \mathbf{p}[i] = \mathbf{p}^\tau[i] && \text{by Def. A.4.} \end{aligned}$$

If $i \in \text{SD}(\mathbf{p})$, then $\mathbf{p}[i] = x_i + p_i(\mathbf{x}, \tau)$ where p_i only depends on variables x_j with $j < i$. For conciseness, define \mathbf{q} to be $\mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau$. Now by idempotence of \mathbf{p} :

$$\begin{aligned}
(\mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau)[i] &= (\mathbf{p} \bullet \mathbf{p} \bullet (\mathbf{p}[1/\tau])^\tau)[i] \\
&= \mathbf{p}[i] \bullet \mathbf{q} \\
&= \mathbf{q}[i] + p_i \bullet \mathbf{q} \\
&\triangleright (\mathbf{p}[1/\tau])^\tau[i] + p_i \bullet \mathbf{q} && \text{since } \mathbf{p}[i] \triangleright x_i \\
&\triangleright x_i + \tau(\mathbf{p}[1/\tau])[i]'' + p_i \bullet \mathbf{q} && \text{by Def. A.4} \\
&\triangleright x_i + \tau\mathbf{p}[i]'' + p_i \bullet \mathbf{q} && \text{since } \mathbf{p}[i]'' \text{ is } \tau\text{-free} \\
&= x_i + p_i \bullet \mathbf{q} + \tau\mathbf{p}[i]'' \\
&\triangleright x_i + p_i \bullet \mathbf{p}^\tau + \tau\mathbf{p}[i]'' && \text{by IH} \\
&\triangleright x_i + p_i \bullet \mathbf{p} + \tau\mathbf{p}[i]'' \\
&= (\mathbf{p} \bullet \mathbf{p})[i] + \tau\mathbf{p}[i]'' \\
&\triangleright \mathbf{p}[i] + \tau\mathbf{p}[i]'' = \mathbf{p}^\tau[i]. \quad \square
\end{aligned}$$

Lemma A.10. *Assume that $\mathbf{p} \in \text{AMPo1}$ is taken from the closure set of a polynomially-bounded SDL. If \mathbf{p} is idempotent, then $(\mathbf{p}) \bullet \mathbf{p} = \mathbf{p}$.*

Proof. $G(\mathbf{p})$ is quasi-acyclic by Lemma A.3 above. Thus, w.l.o.g., it is neat (that is, the variables are indexed to match the DAG order). Clearly, $(\mathbf{p}) \sqsubseteq \mathbf{p}$ and consequently $(\mathbf{p}) \bullet \mathbf{p} \sqsubseteq \mathbf{p} \bullet \mathbf{p} = \mathbf{p}$. Now, assume that equality does not hold. Let \mathbf{m} be a lexicographically biggest monomial⁴ that appears in \mathbf{p} and is missing in $(\mathbf{p}) \bullet \mathbf{p}$. Consider the calculation of $\mathbf{p} \bullet \mathbf{p}$. It generates \mathbf{m} by substituting monomials from the right-hand instance of \mathbf{p} for the variables of a certain monomial \mathbf{n} taken from the left-hand instance. Because \mathbf{p} is neat, substitution replaces a variable x_i either by x_i itself or by a lexicographically-smaller monomial. If all the variables of \mathbf{n} could be replaced by themselves, then $\mathbf{n} = \mathbf{m}$ would be a iterative monomial and therefore present in (\mathbf{p}) and in $(\mathbf{p}) \bullet \mathbf{p}$. We conclude that some variables in \mathbf{n} are not self-dependent. It is further easy to see that \mathbf{m} is lexicographically smaller than \mathbf{n} . By the choice of \mathbf{m} , it follows that \mathbf{n} appears in $(\mathbf{p}) \bullet \mathbf{p}$. But then \mathbf{m} is present in $((\mathbf{p}) \bullet \mathbf{p}) \bullet \mathbf{p}$ which equals $(\mathbf{p}) \bullet \mathbf{p}$ (by associativity of composition and idempotence of \mathbf{p}). A contradiction has been shown. \square

Lemma A.11. *Assume that $\mathbf{p} \in \text{AMPo1}$ is taken from the closure set of a polynomially-bounded SDL. If \mathbf{p} is idempotent, then $(\mathbf{p})^\tau \bullet \mathbf{p} \triangleright \mathbf{p}^\tau$.*

Proof. We already know that every monomial \mathbf{m} of \mathbf{p} appears in $(\mathbf{p}) \bullet \mathbf{p}$ and therefore either \mathbf{m} or $\tau\mathbf{m}$ appears in $(\mathbf{p})^\tau \bullet \mathbf{p}$. So we have covered all the monomials of \mathbf{p}^τ except those affected by generalization, namely those included in the polynomials $\mathbf{p}[i]''$. By definition of $\mathbf{p}[i]''$, they are self-dependent and also appear in (\mathbf{p}) . In fact, $(\mathbf{p})[i] = x_i + \mathbf{p}[i]''$ and $(\mathbf{p})^\tau[i] = x_i + \tau\mathbf{p}[i]''$. Substituting, we find that $((\mathbf{p})^\tau \bullet \mathbf{p})[i] \supseteq x_i + \tau\mathbf{p}[i]''$. The statement of the lemma follows. \square

The proof of Lemma A.7 is now complete, and we move on to the converse subsumption. We will begin this time with an auxiliary result.

⁴Monomials are ordered lexicographically by viewing $x_1^{d_1} \dots x_n^{d_n}$ as $\langle d_n, \dots, d_1 \rangle$. The reversal of the list has the effect that dropping the highest-numbered variable makes a monomial smaller, even if other degrees go up.

Lemma A.12. *Let \mathcal{S} be a polynomially-bounded SDL, and let $\mathbf{p} \in Cl(\mathcal{S})$. Then there is a $k > 0$ such that $\mathbf{p}^{(k)}$, the AMP representation of a k -fold application of \mathbf{p} , is idempotent.*

Proof. This follows from the fact that $Cl(\mathcal{S})$ is finite (an easy consequence of polynomial boundedness). Finiteness implies that the sequence $\mathbf{p}^{(i)}$ must repeat elements. If $\mathbf{p}^{(i)} = \mathbf{p}^{(i+k)}$, $k > 0$, then for any $j \geq i$, $\mathbf{p}^{(j)} = \mathbf{p}^{(j+k)}$. Assuming (w.l.o.g.) $k \geq i$, we have $\mathbf{p}^{(k)} = \mathbf{p}^{(2k)}$, so $\mathbf{p}^{(k)}$ is idempotent. \square

Lemma A.13. *Let \mathcal{S} be a polynomially-bounded SDL, and let \mathbf{q} be an iterative fragment of any $\mathbf{p} \in Cl(\mathcal{S})$. Then*

- (1) $\mathbf{q}^{(i)}$ is also iterative;
- (2) $\mathbf{q}^{(i)} \sqsubseteq \mathbf{p}^{(i)}$;
- (3) $\mathbf{q}^{(i)} \sqsubseteq \mathbf{q}^{(i+1)}$.

Proof. (1) It is easy to see that if $i \in SD(\mathbf{q})$, then $i \in SD(\mathbf{q}^{(i)})$ for any i . On the other hand, as each component of \mathbf{q} only depends on the self-dependent variables, the same is true for $\mathbf{q}^{(i)}$; thus $\mathbf{q}^{(i)}$ is iterative. In fact, it has the same set of self-dependent variables.
(2) This is true because $\mathbf{q} \sqsubseteq \mathbf{p}$ implies $\mathbf{q} \bullet \mathbf{r} \sqsubseteq \mathbf{p} \bullet \mathbf{r}$ for any \mathbf{r} .
(3) Let \mathbf{m} be a monomial of $\mathbf{q}^{(i)}$. When we compose $\mathbf{q}^{(i)}$ over \mathbf{q} (to obtain $\mathbf{q}^{(i+1)}$), we substitute $\mathbf{q}[j]$ for each variable x_j in \mathbf{m} . Since $j \in SD(\mathbf{q})$, we have $x_j \sqsubseteq \mathbf{q}[j]$. Thus, among the monomials resulting from this substitution, there is a monomial identical to \mathbf{m} . \square

Lemma A.14. *Every $\mathbf{p} \in LC(\mathcal{S})$ is subsumed by a member of $SOLVE(\mathcal{S})$.*

Proof. We proceed by three steps, following the definition of LC . First, if $\mathbf{p} \in Cl(\mathcal{S})$ then it is clearly also in $SOLVE(\mathcal{S})$. Next, consider $(\mathbf{p})^\tau$ for $\mathbf{p} \in Cl(\mathcal{S})$. By Lemma A.12, choose k such that $\mathbf{p}^{(k)}$ is idempotent. Hence $(\mathbf{p}^{(k)})^\tau \in SOLVE(\mathcal{S})$. Applying Lemma A.13 to (\mathbf{p}) , we obtain:

$$(\mathbf{p}) \sqsubseteq (\mathbf{p})^{(k)} \sqsubseteq \mathbf{p}^{(k)}$$

Moreover, it is not hard to extend these inequalities to

$$(\mathbf{p})^\tau \sqsubseteq ((\mathbf{p})^{(k)})^\tau \sqsubseteq (\mathbf{p}^{(k)})^\tau$$

Thus $(\mathbf{p})^\tau$ is subsumed by a member of $SOLVE(\mathcal{S})$.

Finally, generalized AMPs are subjected to composition-closure again, but its results will certainly be subsumed by corresponding compositions in $SOLVE(\mathcal{S})$, as the latter is closed under composition, and composition respects the subsumption order. \square

REFERENCES

- [AAG⁺12] Elvira Albert, Puri Arenas, Samir Genaim, German Puebla, and Damiano Zanardini. Cost analysis of object-oriented bytecode programs. *Theoretical Computer Science*, 413(1):142–159, 2012.
- [BA10] Amir M. Ben-Amram. On decidable growth-rate properties of imperative programs. In Patrick Baillot, editor, *International Workshop on Developments in Implicit Computational Complexity (DICE 2010)*, volume 23 of *EPTCS*, pages 1–14, 2010.
- [BAH20] Amir M. Ben-Amram and Geoff W. Hamilton. Tight polynomial worst-case bounds for loop programs. *Logical Methods in Computer Science*, 16(2), 2020.
- [BAK12] Amir M. Ben-Amram and Lars Kristiansen. On the edge of decidability in complexity analysis of loop programs. *International Journal on the Foundations of Computer Science*, 23(7):1451–1464, 2012.

- [BAP16] Amir M. Ben-Amram and Aviad Pineles. Flowchart programs, regular expressions, and decidability of polynomial growth-rate. In Geoff Hamilton, Alexei Lisitsa, and Andrei P. Nemytykh, editors, *Proceedings of the Fourth International Workshop on Verification and Program Transformation (VPT)*, Eindhoven, The Netherlands, volume 216 of *Electronic Proceedings in Theoretical Computer Science*, pages 24–49. Open Publishing Association, 2016.
- [BEF⁺16] Marc Brockschmidt, Fabian Emmes, Stephan Falke, Carsten Fuhs, and Jürgen Giesl. Analyzing runtime and size complexity of integer programs. *ACM Trans. Program. Lang. Syst.*, 38(4):13:1–13:50, August 2016.
- [BJK08] Amir M. Ben-Amram, Neil D. Jones, and Lars Kristiansen. Linear, polynomial or exponential? complexity inference in polynomial time. In Arnold Beckmann, Costas Dimitracopoulos, and Benedikt Löwe, editors, *Logic and Theory of Algorithms, Fourth Conference on Computability in Europe, CiE 2008*, volume 5028 of *LNCS*, pages 67–76. Springer, 2008.
- [Bra06] Mark Braverman. Termination of integer linear programs. In Thomas Ball and Robert B. Jones, editors, *Computer Aided Verification, 18th International Conference, CAV 2006, Seattle, WA, USA*, volume 4144 of *Lecture Notes in Computer Science*, pages 372–385. Springer, 2006.
- [CHS15] Quentin Carbonneaux, Jan Hoffmann, and Zhong Shao. Compositional certified resource bounds. In *Proceedings of the ACM SIGPLAN 2015 Conference on Programming Language Design and Implementation (PLDI)*. ACM, 2015.
- [FH14] Antonio Flores-Montoya and Reiner Hähnle. Resource analysis of complex programs with cost equations. In Jacques Garrigue, editor, *Programming Languages and Systems - 12th Asian Symposium, APLAS 2014, Singapore, Proceedings*, volume 8858 of *Lecture Notes in Computer Science*, pages 275–295. Springer, 2014.
- [GAB⁺17] J. Giesl, C. Aschermann, M. Brockschmidt, F. Emmes, F. Frohn, C. Fuhs, J. Hensel, C. Otto, M. Plücker, P. Schneider-Kamp, T. Ströder, S. Swiderski, and R. Thiemann. Analyzing program termination and complexity automatically with approve. *Journal of Automated Reasoning*, 58(1):3–31, 2017.
- [HFG20] Marcel Hark, Florian Frohn, and Jürgen Giesl. Polynomial loops: Beyond termination. In Elvira Albert and Laura Kovács, editors, *LPAR 2020: 23rd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, Alicante, Spain, May 22-27, 2020*, volume 73 of *EPiC Series in Computing*, pages 279–297. EasyChair, 2020.
- [HOP⁺18] Ehud Hrushovski, Joël Ouaknine, Amaury Pouly, , and James Worrell. Polynomial invariants for affine programs. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '18*, pages 530–539. ACM, 2018.
- [HOW19] Mehran Hosseini, Joël Ouaknine, and James Worrell. Termination of linear loops over the integers. In *46th International Colloquium on Automata, Languages, and Programming, ICALP*, 2019.
- [JK09] Neil D. Jones and Lars Kristiansen. A flow calculus of mwp-bounds for complexity analysis. *ACM Trans. Computational Logic*, 10(4):1–41, 2009.
- [Jon97] Neil D. Jones. *Computability and Complexity From a Programming Perspective*. Foundations of Computing Series. MIT Press, 1997.
- [KA80] Takumi Kasai and Akeo Adachi. A characterization of time complexity by simple loop programs. *Journal of Computer and System Sciences*, 20(1):1–17, 1980.
- [KN04] Lars Kristiansen and Karl-Heinz Niggl. On the computational complexity of imperative programming languages. *Theor. Comp. Sci.*, 318(1-2):139–161, 2004.
- [LDK⁺18] Pedro López-García, Luthfi Darmawan, Maximiliano Klemen, Umer Liqat, Francisco Bueno, and Manuel V. Hermenegildo. Interval-based resource usage verification by translation into horn clauses and an application to energy consumption. *TPLP*, 18(2):167–223, 2018.
- [LM88] Daniel Le Métayer. Ace: an automatic complexity evaluator. *ACM Trans. Program. Lang. Syst.*, 10(2):248–266, 1988.
- [MOS04] Markus Müller-Olm and Helmut Seidl. Computing polynomial program invariants. *Information Processing Letters*, 91(5):233–244, 2004.
- [MR67] Albert R. Meyer and Dennis M. Ritchie. The complexity of loop programs. In *Proc. 22nd ACM National Conference*, pages 465–469. Washington, DC, 1967.
- [NW06] Karl-Heinz Niggl and Henning Wunderlich. Certifying polynomial time and linear/polynomial space for imperative programs. *SIAM J. Comput*, 35(5):1122–1147, 2006.

- [Ros89] M. Rosendahl. Automatic complexity analysis. In *Proceedings of the Conference on Functional Programming Languages and Computer Architecture FPCA '89*, pages 144–156. ACM, 1989.
- [Sch86] Alexander Schrijver. *Theory of Linear and Integer Programming*. John Wiley and Sons, New York, 1986.
- [Ste04] J. Michael Steele. *The Cauchy-Schwarz Master Class: An Introduction to the Art of Mathematical Inequalities*. Cambridge University Press, 2004.
- [SZV17] Moritz Sinn, Florian Zuleger, and Helmut Veith. Complexity and resource bound analysis of imperative programs using difference constraints. *Journal of Automated Reasoning*, 59(1):3–45, 2017.
- [Tiw04] Ashish Tiwari. Termination of linear programs. In *Computer Aided Verification*, pages 70–82, 2004.
- [Weg75] Ben Wegbreit. Mechanical program analysis. *Communications of the ACM*, 18(9):528–539, 1975.