

Complexity of Bradley-Manna-Sipma Lexicographic Ranking Functions

Amir M. Ben-Amram¹ and Samir Genaim²(✉)

¹ School of Computer Science, The Tel-Aviv Academic College, Tel Aviv, Israel

² DSIC, Complutense University of Madrid (UCM), Madrid, Spain
genaim@gmail.com

Abstract. In this paper we turn the spotlight on a class of lexicographic ranking functions introduced by Bradley, Manna and Sipma in a seminal CAV 2005 paper, and establish for the first time the complexity of some problems involving the inference of such functions for linear-constraint loops (without precondition). We show that finding such a function, if one exists, can be done in polynomial time in a way which is sound and complete when the variables range over the rationals (or reals). We show that when variables range over the integers, the problem is harder—deciding the existence of a ranking function is coNP-complete. Next, we study the problem of minimizing the number of components in the ranking function (a.k.a. the dimension). This number is interesting in contexts like computing iteration bounds and loop parallelization. Surprisingly, and unlike the situation for some other classes of lexicographic ranking functions, we find that even deciding whether a two-component ranking function exists is harder than the unrestricted problem: NP-complete over the rationals and Σ_2^P -complete over the integers.

1 Introduction

Proving that a program will not go into an infinite loop is one of the most fundamental tasks of program verification, and has been the subject of voluminous research. Perhaps the best known, and often used, technique for proving termination is the *ranking function*. This is a function ρ that maps the program states into the elements of a well-founded ordered set, such that $\rho(s) > \rho(s')$ holds for any consecutive states s and s' . This implies termination since infinite descent in a well-founded order is impossible.

We focus on *numerical loops*, where a state is described by the values of a finite set of numerical variables; we consider the setting of integer-valued variables, as well as rational-valued (or real-valued) variables. We ignore details of the programming language; we assume that we are provided an abstract description of the loop as a finite number of alternatives, that we call *paths*, each one

This work was funded partially by the EU project FP7-ICT-610582 ENVISAGE: Engineering Virtualized Services (<http://www.envisage-project.eu>), by the Spanish MINECO project TIN2012-38137, and by the CM project S2013/ICE-3006.

defined by a finite set of *linear constraints* on the program variables x, y, \dots and the primed variables x', y', \dots which refer to the state following the iteration. The following is such a loop consisting of four paths, $\mathcal{Q}_1, \dots, \mathcal{Q}_4$:

$$\begin{aligned} \mathcal{Q}_1 &= \{x \geq 0, x' \leq x - 1, & y' &= y, & z' &= z\} \\ \mathcal{Q}_2 &= \{x \geq 0, x' \leq x - 1, & y' &= y, & z \geq 0, z' \leq z - 1\} \\ \mathcal{Q}_3 &= \{ & x' &= x, & y \geq 0, y' \leq y - 1, z \geq 0, z' \leq z - 1\} \\ \mathcal{Q}_4 &= \{ & x' &= x, & y \geq 0, y' \leq y - 1, & z' &= z\} \end{aligned}$$

Note that \mathcal{Q}_i are convex polyhedra. A transition from a state \bar{x} to \bar{x}' is possible iff (\bar{x}, \bar{x}') is a point in some path \mathcal{Q}_i . We remark that our results hold for arbitrarily-complex control-flow graphs (CFGs), we prefer to use the loop setting for clarity.

A popular tool for proving the termination of such loops is *linear ranking functions* (LRFs). An LRF is a function $\rho(x_1, \dots, x_n) = a_1x_1 + \dots + a_nx_n + a_0$ such that any transition (\bar{x}, \bar{x}') satisfies (i) $\rho(\bar{x}) \geq 0$; and (ii) $\rho(\bar{x}) - \rho(\bar{x}') \geq 1$. E.g., $\rho(x, y, z) = x$ is an LRF for a loop that consists of only \mathcal{Q}_1 and \mathcal{Q}_2 above, $\rho(x, y, z) = y$ is an LRF for \mathcal{Q}_3 and \mathcal{Q}_4 , and $\rho(x, y, z) = z$ is an LRF for \mathcal{Q}_2 and \mathcal{Q}_3 . However, there is no LRF that satisfies the above conditions for all paths $\mathcal{Q}_1, \dots, \mathcal{Q}_4$. An algorithm to find an LRF using linear programming (LP) has been found by multiple researchers in different places and times and in some alternative versions [1, 10, 14, 21, 23, 26]. Since LP has a polynomial-time complexity, most of these methods yield polynomial-time algorithms. These algorithms are complete for loops with rational-valued variables, but not with integer-valued variables. Indeed, [4] shows loops that have LRFs over the integers but do not even terminate over the rationals. In a previous work [4] we considered the integer setting, where complete algorithms were proposed and a complexity classification was proved: to decide whether an LRF exists is *coNP-complete*.

LRFs do not suffice for all loops (e.g., the 4-path loop above), and thus, a natural question is what to do when an LRF does not exist; and a natural answer is to try a richer class of ranking functions. Of particular importance is the class of *lexicographic-linear ranking functions* (LLRFs). An LLRF is a d -tuple of affine-linear functions, $\langle \rho_1, \dots, \rho_d \rangle$, required to descend lexicographically. Interestingly, Alan Turing's early demonstration [28] of how to verify a program used an LLRF for the termination proof. *Algorithms* to find LLRFs for linear-constraint loops (or CFGs) can use LP techniques, extending the work on LRFs. Alias et al. [1] extended the polynomial-time LRF algorithm to LLRFs and gave a complete solution for CFGs. As for LRFs, the solution is incomplete for integer data, and in [4] we established for LLRFs over the integers results that parallel those for LRFs, in particular, to decide whether an LLRF exists is *coNP-complete*.

Interestingly, when trying to define the requirements from a numeric “lexicographic ranking function” (corresponding to the conditions (i) and (ii) on an LRF, above), different researchers had come up with different definitions. In particular, the definition in [1] is more restrictive than the definition in [4]. Furthermore, an important paper [5] on LLRF generation that preceded both works gave yet a different definition. We give the precise definitions in Sect. 2; for the purpose of introduction, let us focus on the LLRFs of [5] (henceforth, BMS-LLRFs, after the authors), and illustrate the definition by an example.

Consider the above loop defined by $\mathcal{Q}_1, \dots, \mathcal{Q}_4$. A possible BMS-LLRF for this loop is $\rho(x, y, z) = \langle x, y \rangle$. The justification is this: in \mathcal{Q}_1 and \mathcal{Q}_2 , the function $\rho_1(x, y) = x$ is ranking (non-negative and decreasing by at least 1). In \mathcal{Q}_3 and \mathcal{Q}_4 , $\rho_2(x, y) = y$ is ranking, while ρ_1 is non-increasing. This is true over the rationals and *a fortiori* over the integers. The following points are important: (1) for each path we have an LRF, which is one of the components of the BMS-LLRF; and (2) previous (lower-numbered) components are only required to be non-increasing on that path. Note that this LLRF does not satisfy the requirements of [1] or [4].

The goal of this paper is to understand the *computational complexity* of some problems related to BMS-LLRFs, starting with the most basic problem, whether a given loop has such LLRF. We note that [5] does not provide an answer, as a consequence of attempting to solve a much harder problem—they consider a loop given with a precondition and search for a BMS-LLRF together with a supporting linear invariant. We do not know if this problem is even decidable when parameters like the number of constraints in the invariants are not fixed in advance (when they are, the approach of [5] is complete, but only over the reals, and at a high computational cost – even without a precondition).

We consider the complexity of finding a BMS-LLRF for a given loop, without preconditions. We prove that this can be done in polynomial time when the loop is interpreted over the rationals, while over the integers, deciding the existence of a BMS-LLRF is coNP-complete. An exponential-time synthesis algorithm is also given. These results are similar to those obtained for the previously studied classes of LLRFs [4], but are shown for the first time for BMS-LLRFs.

Next, we consider the number of components d in a BMS-LLRF $\langle \rho_1, \dots, \rho_d \rangle$. This number is informally called the *dimension* of the function. It is interesting for several reasons: An upper bound on the dimension is useful for fixing the template in the constraint-solving approach, and plays a role in analyzing the complexity of corresponding algorithms. In addition, an LLRF can be used to infer bounds on the number of iterations [1]; assuming linear bounds on individual variables, a polynomial bound of degree d is clearly implied, which motivates the desire to minimize the dimension, to obtain tight bounds. A smaller dimension also means better results when LLRFs are used to guide parallelization [15].

Importantly, the algorithms of Alias et al. [1] and Ben-Amram and Genaim [4] are optimal w.r.t. the dimension, i.e., they synthesize LLRFs of minimal dimension for the respective classes. We note that it is possible for a loop to have LLRFs of all three classes but such that the minimal dimension is different in each (see Sect. 4). We also note that, unlike the case for the previous classes, our synthesis algorithm for BMS-LLRFs is *not* guaranteed to produce a function of minimal dimension. This leads us to ask: (1) what is the best *a priori* bound on the dimension, in terms of the number of variables and paths; and (2) how difficult it is to find an LLRF of minimal dimension. As a relaxation of this optimization problem, we can pose the problem of finding an LLRF that satisfies a given bound on the dimension. Our results are summarized in Table 1. There is a striking difference of BMS-LLRFs from other classes w.r.t. to the minimum dimension problem: the complexity jumps from PTIME (resp. coNP-complete)

to NPC (resp. Σ_2^P -complete) over rationals (resp. integers). This holds for any fixed dimension larger than one (dimension one is an LRF).

Table 1. Summary of results, considering a loop of k paths over n variables. Those in the third row are new, the others are from previous works or follow by minor variations.

LLRF type	Dimension bound	Existence		Fixed dimension	
		over \mathbb{Q}	over \mathbb{Z}	over \mathbb{Q}	over \mathbb{Z}
ADFG [1]	$\min(n, k)$	PTIME	coNP-complete	PTIME	coNP-complete
BG [4]	n	PTIME	coNP-complete	PTIME	coNP-complete
BMS [5]	k	PTIME	coNP-complete	NP-complete	Σ_2^P -complete

2 Preliminaries

Polyhedra. A rational convex polyhedron $\mathcal{P} \subseteq \mathbb{Q}^n$ (polyhedron for short) is 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. We say that \mathcal{P} is specified by $A\mathbf{x} \leq \mathbf{b}$. We use calligraphic letters, such as \mathcal{P} and \mathcal{Q} to denote polyhedra. For a given polyhedron $\mathcal{P} \subseteq \mathbb{Q}^n$ we let $I(\mathcal{P})$ be $\mathcal{P} \cap \mathbb{Z}^n$, i.e., the set of integer points of \mathcal{P} . The integer hull of \mathcal{P} , commonly denoted by \mathcal{P}_I , is defined as the convex hull of $I(\mathcal{P})$. It is known that \mathcal{P}_I is also a polyhedron. An integer polyhedron is a polyhedron \mathcal{P} such that $\mathcal{P} = \mathcal{P}_I$. We also say that \mathcal{P} is integral.

Multipath Linear-Constraint Loops. A multipath linear-constraint loop (MLC loop) with k paths has the form: $\bigvee_{i=1}^k A_i \begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix} \leq \mathbf{c}_i$ where $\mathbf{x} = (x_1, \dots, x_n)^T$ and $\mathbf{x}' = (x'_1, \dots, x'_n)^T$ are column vectors, and for $q > 0$, $A_i \in \mathbb{Q}^{q \times 2n}$, $\mathbf{c}_i \in \mathbb{Q}^q$. Each path $A_i \begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix} \leq \mathbf{c}_i$ is called an abstract transition. The loop is a rational loop if \mathbf{x} and \mathbf{x}' range over \mathbb{Q}^n , and it is an integer loop if they range over \mathbb{Z}^n . We say that there is a transition from a state $\mathbf{x} \in \mathbb{Q}^n$ to a state $\mathbf{x}' \in \mathbb{Q}^n$, if for some $1 \leq i \leq k$, $\begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix}$ satisfies the i -th abstract transition. In such case we say that \mathbf{x} is an enabled state. We use \mathbf{x}'' as a shorthand for a transition $\begin{pmatrix} \mathbf{x} \\ \mathbf{x}' \end{pmatrix}$, and consider it as a point in \mathbb{Q}^{2n} . The set of transitions satisfying a particular abstract transition is a polyhedron in \mathbb{Q}^{2n} , denoted \mathcal{Q}_i , namely $A_i \mathbf{x}'' \leq \mathbf{c}_i$. In our work it is convenient to represent an MLC loop by its transition polyhedra $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, which we often write with explicit equalities and inequalities. These are sometimes referred to as the paths of the multipath loop.

Ranking Functions. An affine linear function $\rho : \mathbb{Q}^n \mapsto \mathbb{Q}$ is of the form $\rho(\mathbf{x}) = \boldsymbol{\lambda} \cdot \mathbf{x} + \lambda_0$ where $\boldsymbol{\lambda} \in \mathbb{Q}^n$ and $\lambda_0 \in \mathbb{Q}$. We define $\Delta\rho : \mathbb{Q}^{2n} \mapsto \mathbb{Q}$ as $\Delta\rho(\mathbf{x}'') = \rho(\mathbf{x}) - \rho(\mathbf{x}')$. Given a set $T \subseteq \mathbb{Q}^{2n}$, representing transitions, we say that ρ is an LRF for T if for every $\mathbf{x}'' \in T$ we have (i) $\rho(\mathbf{x}) \geq 0$; and (ii) $\Delta\rho(\mathbf{x}'') \geq 1$. We say that ρ is an LRF for a rational (resp. integer) loop,

specified by $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, when it is an LRF for $\bigcup_{i=1}^k \mathcal{Q}_i$ (resp. $\bigcup_{i=1}^k I(\mathcal{Q}_i)$). For a rational loop, there is a polynomial-time algorithm to either find an LRF or determine that none exists [23]. Its essence is that using *Farkas' Lemma* [25, p. 93], it is possible to set up an LP problem whose feasibility is equivalent to the existence of ρ that satisfies (i) and (ii) over $\mathcal{Q}_1, \dots, \mathcal{Q}_k$.

A *d-dimensional affine function* $\tau : \mathbb{Q}^n \rightarrow \mathbb{Q}^d$ is expressed by a d -tuple $\tau = \langle \rho_1, \dots, \rho_d \rangle$, where each component $\rho_i : \mathbb{Q}^n \rightarrow \mathbb{Q}$ is an affine linear function. The number d is informally called the *dimension* of τ . Next we define when such a function is BMS-LLRF [5] for a given rational or integer MLC loop. We then compare with ADFG-LLRFs (due to [1]) and BG-LLRFs (due to [4]).

Definition 1 (BMS-LLRF). *Given k sets of transitions $T_1, \dots, T_k \subseteq \mathbb{Q}^{2n}$, we say that $\tau = \langle \rho_1, \dots, \rho_d \rangle$ is a BMS-LLRF for T_1, \dots, T_k iff for every $1 \leq \ell \leq k$ there is $1 \leq i \leq d$ such that the following hold for any $\mathbf{x}'' \in T_\ell$:*

$$\forall j < i . \Delta\rho_j(\mathbf{x}'') \geq 0, \tag{1}$$

$$\rho_i(\mathbf{x}) \geq 0, \tag{2}$$

$$\Delta\rho_i(\mathbf{x}'') \geq 1. \tag{3}$$

We say that T_ℓ is ranked by ρ_i .

We say that τ is a BMS-LLRF for a rational (resp. integer) loop, specified by $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, when it is a BMS-LLRF for $\mathcal{Q}_1, \dots, \mathcal{Q}_k$ (resp. $I(\mathcal{Q}_1), \dots, I(\mathcal{Q}_k)$). It is easy to see that the existence of a BMS-LLRF implies termination.

Definition 2 (BG-LLRF). *Given a set of transitions $T \subseteq \mathbb{Q}^{2n}$, we say that $\tau = \langle \rho_1, \dots, \rho_d \rangle$ is a BG-LLRF for T iff for every $\mathbf{x}'' \in T$ there is $1 \leq i \leq d$ such that the following hold:*

$$\forall j < i . \Delta\rho_j(\mathbf{x}'') \geq 0, \tag{4}$$

$$\forall j \leq i . \rho_j(\mathbf{x}) \geq 0, \tag{5}$$

$$\Delta\rho_i(\mathbf{x}'') \geq 1. \tag{6}$$

We say that \mathbf{x} is ranked by ρ_i .

We say that τ is a BG-LLRF for a rational (resp. integer) loop, specified by $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, when it is a BG-LLRF for $\mathcal{Q}_1 \cup \dots \cup \mathcal{Q}_k$ (resp. $I(\mathcal{Q}_1) \cup \dots \cup I(\mathcal{Q}_k)$). It is easy to see that the existence of a BG-LLRF implies termination.

Note the differences between the definitions: in one sense, BG-LLRFs are more flexible because of the different quantification — for every transition \mathbf{x}'' there has to be a component ρ_i that ranks it, but i may differ for different \mathbf{x}'' , whereas in BMS-LLRFs, all transitions that belong to a certain T_ℓ have to be ranked by the same component. In another sense, BMS-LLRFs are more flexible because components ρ_j with $j < i$ can be negative (compare (2) with (5)). Thus, there are loops that have a BMS-LLRF and do not have a BG-LLRF (see loop in Sect. 1); and vice versa (see [4, Ex. 2.12]). A third type of LLRFs is attributed to [1], hence we refer to it as ADFG-LLRF. It is similar to BG-LLRFs

but requires all components to be non-negative in every enabled state. That is, condition (5) is strengthened. Interestingly, the completeness proof in [1] shows that the above-mentioned flexibility of BG-LLRFs adds no power in this case; therefore, ADFG-LLRFs are a special case of both BG-LLRFs and BMS-LLRFs.

The decision problem *Existence of a BMS-LLRF* deals with deciding whether a given MLC loop admits a BMS-LLRF, we denote it by $\text{BMS-LEXLINRF}(\mathbb{Q})$ and $\text{BMS-LEXLINRF}(\mathbb{Z})$ for rational and integer loops respectively. The corresponding decision problems for ADFG- and BG-LLRFs are solved in [1] and [4], respectively, over the rationals; the case of integers is only addressed in [4] for BG-LLRFs, but the complexity results apply to ADFG-LLRFs as well.

3 Synthesis of BMS-LLRFs

In this section we describe a complete algorithm for synthesizing BMS-LLRFs for rational and integer MLC loops; and show that the decision problems $\text{BMS-LEXLINRF}(\mathbb{Q})$ and $\text{BMS-LEXLINRF}(\mathbb{Z})$ are PTIME and coNP-complete, respectively. We assume a given MLC loop $\mathcal{Q}_1, \dots, \mathcal{Q}_k$ where each \mathcal{Q}_i is given as a set of linear constraints, over $2n$ variables (n variables and n primed variables).

Definition 3. Let T_1, \dots, T_k be sets of transitions such that $T_i \subseteq \mathbb{Q}^{2n}$. We say that an affine linear function ρ is a BMS quasi-LRF (BMS-QLRF for short) for T_1, \dots, T_k if every transition $\mathbf{x}'' \in T_1 \cup \dots \cup T_k$ satisfies $\Delta\rho(\mathbf{x}'') \geq 0$, and for at least one T_ℓ , ρ is an LRF (such T_ℓ is said to be ranked by ρ).

Example 1. The following are BMS-QLRFs for the loop consisting of $\mathcal{Q}_1, \dots, \mathcal{Q}_4$ presented in Sect. 1: $f_1(x, y, z) = x$, which ranks $\{\mathcal{Q}_1, \mathcal{Q}_2\}$; $f_2(x, y, z) = y$ which ranks $\{\mathcal{Q}_3, \mathcal{Q}_4\}$; and $f_3(x, y, z) = z$ which ranks $\{\mathcal{Q}_2, \mathcal{Q}_3\}$.

Lemma 1. *There is a polynomial-time algorithm that finds a BMS-QLRF ρ , if there is any, for $\mathcal{Q}_1, \dots, \mathcal{Q}_k$.*

Proof. The algorithm iterates over the paths $\mathcal{Q}_1, \dots, \mathcal{Q}_k$. In the i -th iteration it checks if there is an LRF ρ for \mathcal{Q}_i that is non-increasing for all other paths, stopping if it finds one. The algorithm makes at most k iterations. Each iteration can be implemented in polynomial time using Farkas' Lemma (as in [23]). \square

Our procedure for synthesizing BMS-LLRFs is depicted in Algorithm 1. In each iteration (i.e., call to LLRFSYN): it finds a BMS-QLRF ρ for the current paths (Line 2); it eliminates all paths that are ranked by ρ (Line 3); and calls recursively to handle the remaining paths (Line 4). The algorithm stops when all paths are ranked (Line 1), or when it does not find a BMS-QLRF (Line 6).

Example 2. Consider the MLC loop example in Sect. 1. Procedure LLRFSYN is first applied to $\langle \mathcal{Q}_1, \mathcal{Q}_2, \mathcal{Q}_3, \mathcal{Q}_4 \rangle$, and at Line 2 we can choose the BMS-QLRF x which ranks \mathcal{Q}_1 and \mathcal{Q}_2 . Hence these are eliminated at Line 3, and at Line 4 LLRFSYN is applied recursively to $\langle \emptyset, \emptyset, \mathcal{Q}_3, \mathcal{Q}_4 \rangle$. Then at Line 2 we can choose the BMS-QLRF y which ranks \mathcal{Q}_3 and \mathcal{Q}_4 . The next recursive call receives empty polyhedra, and thus the check at Line 1 succeeds and the algorithm returns $\langle x, y \rangle$.

Algorithm 1. Synthesizing BMS-LLRFs

```

LLRFSYN( $\langle \mathcal{Q}_1, \dots, \mathcal{Q}_k \rangle$ )
begin
1  if  $\langle \mathcal{Q}_1, \dots, \mathcal{Q}_k \rangle$  are all empty then return nil
2  if  $\mathcal{Q}_1, \dots, \mathcal{Q}_k$  has a BMS-QLRF  $\rho$  then
3       $\forall 1 \leq i \leq k. \mathcal{Q}'_i := \emptyset$  if  $\mathcal{Q}_i$  is ranked by  $\rho$ , otherwise  $\mathcal{Q}'_i = \mathcal{Q}_i$ 
4       $\tau \leftarrow$  LLRFSYN( $\langle \mathcal{Q}'_1, \dots, \mathcal{Q}'_k \rangle$ )
5      if  $\tau \neq \text{NONE}$  then return  $\rho :: \tau$ 
6  return NONE

```

Lemma 2. *If LLRFSYN($\langle \mathcal{Q}_1, \dots, \mathcal{Q}_k \rangle$) returns τ different from NONE, then τ is a BMS-LLRF for the rational loop $\mathcal{Q}_1, \dots, \mathcal{Q}_k$.*

The proof of the above lemma is straightforward. Thus, Algorithm 1 is a sound algorithm for BMS-LLRFs. The following proposition shows completeness.

Proposition 1. *There is a BMS-LLRF for $\mathcal{Q}_1, \dots, \mathcal{Q}_k$ if and only if every subset of $\{\mathcal{Q}_1, \dots, \mathcal{Q}_k\}$ has a BMS-QLRF.*

Proof. The “if” direction is implied by the LLRFSYN procedure, in such case it will find a BMS-LLRF. For the “only if” direction, let $\tau = \langle \rho_1, \dots, \rho_d \rangle$ be a BMS-LLRF for $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, and let $\mathcal{Q}_{\ell_1}, \dots, \mathcal{Q}_{\ell_j}$ be an arbitrary subset of the loop’s paths. Since τ is a BMS-LLRF for $\mathcal{Q}_1, \dots, \mathcal{Q}_k$, each \mathcal{Q}_{ℓ_i} is ranked by some ρ_{l_i} . Let $l = \min\{l_1, \dots, l_j\}$, then ρ_l is a BMS-QLRF for $\mathcal{Q}_{\ell_1}, \dots, \mathcal{Q}_{\ell_j}$. \square

Lemma 3. *Procedure LLRFSYN can be implemented in polynomial time.*

Proof. Procedure LLRFSYN makes at most k steps (since at least one path is eliminated in every step). Further, all steps are elementary except checking for a BMS-QLRF which can be done in polynomial time as stated by Lemma 1. \square

Corollary 1. $\text{BMS-LEXLINRF}(\mathbb{Q}) \in \text{PTIME}$.

So far we have considered only rational loops, next we consider integer loops.

Lemma 4. *There is a complete algorithm for synthesizing a BMS-QLRF for $I(\mathcal{Q}_1), \dots, I(\mathcal{Q}_k)$.*

Proof. The algorithm computes the integer hull $\mathcal{Q}_{1I}, \dots, \mathcal{Q}_{kI}$, and then proceeds as in the rational case (Lemma 1). Correctness follows from the fact that for integral polyhedra the implied inequalities over the rationals and integers coincide, i.e., $\mathcal{Q}_{1I}, \dots, \mathcal{Q}_{kI}$ and $I(\mathcal{Q}_1), \dots, I(\mathcal{Q}_k)$ have the same BMS-QLRFs. \square

Lemma 5. *When procedure LLRFSYN is applied to the integer hulls $\mathcal{Q}_{1I}, \dots, \mathcal{Q}_{kI}$, it finds a BMS-LLRF for $I(\mathcal{Q}_1), \dots, I(\mathcal{Q}_k)$, if one exists.*

Proof. Soundness follows from the fact that \mathcal{Q}_I contains $I(\mathcal{Q})$; for completeness, note that: (i) Proposition 1 holds also for integer loops; and (ii) Line 3 of LLRFSYN does not change the transition polyhedra, it only eliminates some, which means that they remain integral throughout the recursive calls. Thus, in each iteration the check at Line 2 is complete (see Lemma 4). \square

In the general case this procedure has an exponential time complexity since computing the integer hull requires an exponential time. However, for special cases in which the integer hull can be computed in polynomial time [4, Sect. 4] it has polynomial time complexity. The following lemma implies (assuming $P \neq NP$) that the exponential time complexity is unavoidable in general.

Theorem 1. *BMS-LEXLINRF(\mathbb{Z}) is a coNP-complete problem.*

Proof. The coNP-hardness follows from the reduction in [4, Sect. 3.1], since it constructs a loop that either does not terminate or has an LRF. The inclusion in coNP is based on arguments similar to those in [4, Sect. 5]; briefly, we use the generator representation of the transition polyhedra to construct a polynomial-size witness against existence of an LLRF (see [2]). \square

4 The Dimension of BMS-LLRFs

Ben-Amram and Genaim [4, Cor. 5.12, p. 32] showed that if a given MLC loop has a BG-LLRF, then it has one of dimension at most n , the dimension of the state space. The same proof can be used to bound the dimension of ADFG-LLRFs by n as well. Hence for ADFG-LLRFs the bound $\min(n, k)$ holds (k is the number of paths), due to the fact that associating LLRF components with paths is no loss of generality for ADFG-LLRFs [1]. In the case of BMS-LLRFs, the bound k clearly holds, and the next example shows that it is tight.

Example 3. Define an MLC loop $\mathcal{Q}_1, \dots, \mathcal{Q}_k$ for some $k > 0$, over variables x, y , where each $\mathcal{Q}_i = \{x' \leq x, x' + i \cdot y' \leq x + i \cdot y - 1, x + i \cdot y \geq 0\}$. Define $f_i(x, y) = x + i \cdot y$. It is easy to check that (i) f_i is an LRF for \mathcal{Q}_i , and is non-increasing for any \mathcal{Q}_j with $i < j \leq k$; and (ii) there are no distinct \mathcal{Q}_i and \mathcal{Q}_j that have a common LRF. From (i) it follows that $\langle f_1, \dots, f_k \rangle$ is a BMS-LLRF for this loop, and from (ii) it follows that any BMS-LLRF must have (at least) dimension k , since different paths cannot be ranked by the same component. We remark that this loop has no BG-LLRF (hence, also no ADFG-LLRF).

The above discussion emphasizes the difference between the various definitions of LLRFs, when considering the dimension. The next example emphasizes this difference further, it shows that there are loops, having LLRFs of all three kinds, for which the minimal dimension is different according to each definition. This also means that the implied bounds on the number of iterations (assuming, for simplicity, that all variables have the same upper bound) are different.

Example 4. Consider an MLC loop specified by the following paths

$$\begin{aligned}
 \mathcal{Q}_1 &= \left\{ \begin{array}{llll} r \geq 0, & t \geq 0, & x \geq 0, & z \geq 0, \quad w \geq 0, \\ r' < r, & t' < t, & & \end{array} \right\} \\
 \mathcal{Q}_2 &= \left\{ \begin{array}{llll} r \geq 0, & s \geq 0, & t \geq 0, & x \geq 0, \quad z \geq 0, \quad w \geq 0, \\ r' = r, & s' < s, & t' < t, & \end{array} \right\} \\
 \mathcal{Q}_3 &= \left\{ \begin{array}{llll} r \geq 0, & s \geq 0, & t' = t & x \geq 0, \quad z \geq 0, \quad w \geq 0, \\ r' = r, & s' = s, & & x' < x, \end{array} \right\} \\
 \mathcal{Q}_4 &= \left\{ \begin{array}{llll} r \geq 0, & s \geq 0, & t' = t & x \geq 0, \quad y \geq 0, \quad z \geq 0, \quad w \geq 0, \\ r' = r, & s' = s, & & x' = x, \quad y' < y, \quad z' < z, \end{array} \right\} \\
 \mathcal{Q}_5 &= \left\{ \begin{array}{llll} r \geq 0, & s \geq 0, & t' = t & x \geq 0, \quad y \geq 0, \quad z \geq 0, \quad w \geq 0, \\ r' = r, & s' = s, & & x' = x, \quad y' < y, \quad z' = z, \quad w' < w \end{array} \right\}
 \end{aligned}$$

where, for readability, we use $<$ for the relation “smaller at least by 1”. This loop has the BMS-LLRF $\langle t, x, y \rangle$, which is neither a BG-LLRF or ADFG-LLRF because t is not lower-bounded on all the paths. Its shortest BG-LLRF is of dimension 4, e.g., $\langle r, s, x, y \rangle$, which is not an ADFG-LLRF because y is not lower-bounded on all the paths. Its shortest ADFG-LLRF is of dimension 5, e.g., $\langle r, s, x, z, w \rangle$. This reasoning is valid for both integer and rational variables.

Next, we consider the problem of minimal dimension. We ask (1) whether our algorithms return an LLRF with minimal dimension; and (2) what do we gain (or lose?) in terms of computational tractability if we fix a bound on the dimension in advance. Importantly, the algorithms of [1, 4] are optimal w.r.t. the dimension, i.e., they synthesize LLRFs of minimal dimension. In both cases the optimal result is obtained by a greedy algorithm, that constructs the LLRF by adding one dimension at a time, taking care in each iteration to rank as many transitions as possible. The next example shows that a greedy choice in Algorithm 1 fails to guarantee optimality, for both rational and integer loops. Intuitively, the greedy approach worked in [1, 4] because the classes of quasi-LLRFs used to construct LLRFs are closed under conic combinations, so there is always an optimal choice that dominates all others. This is not true for BMS-QLRFs.

Example 5. Consider the MLC loop of Sect. 1. If at Line 2 of Algorithm 1 we seek a BMS-QLRF that ranks a maximal number of the paths, we can use any of those derived in Example 1: $f_1 = x$; $f_2 = y$; or $f_3 = z$. However, these alternatives lead to BMS-LLRFs of different dimensions: (i) choose f_1 to rank $\{\mathcal{Q}_1, \mathcal{Q}_2\}$, and then f_2 to rank $\{\mathcal{Q}_3, \mathcal{Q}_4\}$. (ii) choose f_2 to rank $\{\mathcal{Q}_3, \mathcal{Q}_4\}$, and then f_1 to rank $\{\mathcal{Q}_1, \mathcal{Q}_2\}$. (iii) choose f_3 to rank $\{\mathcal{Q}_2, \mathcal{Q}_3\}$, but then there is no single function that ranks $\{\mathcal{Q}_1, \mathcal{Q}_4\}$. Take f_1 to rank \mathcal{Q}_1 and then f_2 to rank \mathcal{Q}_4 . The dimension of the BMS-LLRF in the first two cases is 2, and in the last one it is 3.

Since Algorithm 1 is not guaranteed to find a BMS-LLRF of minimal dimension, it is natural to ask *how hard is the problem of finding a BMS-LLRF of minimal dimension?* This can be posed as a decision problem: *does a given MLC loop have a BMS-LLRF with dimension at most d ?* This decision problem is denoted by $\text{BMS-LEXLINRF}(d, \mathbb{Q})$ and $\text{BMS-LEXLINRF}(d, \mathbb{Z})$ for rational and integer loops respectively. Note that d is a constant, however, it will be clear that accepting

d as an input does not change the complexity class of these problems. Also note that for $d = 1$ it is just the LRF problem. Similar problems can be formulated for ADFG- and BG-LLRFs, of course. In these two settings, the imposition of a dimension bound does not change the complexity class.

Theorem 2. *Given a rational MLC loop, and $d \geq 1$, it is possible to determine in polynomial time if there is an ADFG-LLRF (resp. BG-LLRFs) for the loop of dimension at most d . For integer MLC loops, the problem is coNP-complete.*

Proof. The case of rational loops is straightforward since the corresponding synthesis algorithms find LLRFs with minimal dimension, and are in PTIME. The integer case follows easily from the techniques of [4] (see [2]). \square

5 Complexity of BMS-LEXLINRF (d, \mathbb{Q})

In this section we show that BMS-LEXLINRF(d, \mathbb{Q}) is NP-complete.

Theorem 3. *For $d \geq 2$, BMS-LEXLINRF(d, \mathbb{Q}) is an NP-complete problem.*

For inclusion in NP, a non-deterministic algorithm for the problem works as follows. First, it *guesses* a partition of $\{1, \dots, k\}$ into d sets J_1, \dots, J_d , of which some may be empty (we can assume they are last). Then it proceeds as in Algorithm 1 but insists that the paths indexed by J_r be ranked at the r -th iteration. This may fail, and then the algorithm rejects. If a BMS-LLRF of dimension at most d exists, there will be an accepting computation.

For NP-hardness we reduce from the NP-complete problem *d-Colorability of 3-Uniform Hypergraphs* [20, 22]. An instance of this problem is a set H of m sets F_1, \dots, F_m (hyperedges, or “faces”), where each F_i includes exactly 3 elements from a set of vertices $V = \{1, \dots, n\}$, and we are asked whether we can choose a color (out of d colors) for each vertex such that every face is not monocolored.

We construct a rational MLC loop in $3m$ variables and n paths. The variables are indexed by vertices and faces: variable $x_{i,j}$ corresponds to $i \in F_j \in H$. For each vertex $1 \leq i \leq n$ we define Q_i as a conjunction of the following:

$$\sum_{k: i \in F_k} x_{i,k} - \sum_{k: i \in F_k} x'_{i,k} \geq 1 \tag{7}$$

$$\sum_{k: j \in F_k} x_{j,k} - \sum_{k: j \in F_k} x'_{j,k} \geq 0 \quad \text{for all vertex } j \neq i \tag{8}$$

$$x_{i,k} \geq 0 \quad \text{for all face } F_k \text{ s.t. } i \in F_k \tag{9}$$

$$x_{j,k} \geq 0 \quad \text{for all vertex } j \text{ and face } F_k \text{ s.t. } j \in F_k \wedge i \notin F_k \tag{10}$$

$$x_{i,k} + x_{j,k} \geq 0 \quad \text{for all vertex } j \neq i \text{ and face } F_k \text{ s.t. } i, j \in F_k \tag{11}$$

We claim that a rational loop that consists of these n paths has a BMS-LLRF of dimension d iff there is a valid d -coloring for the vertices V .

Assume given a d -coloring, namely a division of the vertices in d disjoint sets $V = C_1 \cup \dots \cup C_d$, such that the vertices of each C_i are assigned the same color. We construct a BMS-LLRF $\langle g_1, \dots, g_d \rangle$ such that g_ℓ ranks all paths Q_i with $i \in C_\ell$. We assume that each C_ℓ is non-empty (otherwise we let $g_\ell(\mathbf{x}) = 0$).

We start with C_1 . For each $F_k \in H$, define a function f_k as follows: if $F_k \cap C_1 = \emptyset$ we let $f_k(\mathbf{x}) = 0$; if $F_k \cap C_1 = \{i\}$ we let $f_k(\mathbf{x}) = x_{i,k}$; and if $F_k \cap C_1 = \{i, j\}$ we let $f_k(\mathbf{x}) = x_{i,k} + x_{j,k}$. We claim that $g_1(\mathbf{x}) = \sum_k f_k$ is a BMS-QLRF for Q_1, \dots, Q_n that ranks all paths Q_i with $i \in C_1$, which we justify as follows:

1. g_1 is non-increasing on all Q_j , and decreasing for each Q_i with $i \in C_1$. To see this, rewrite $g(\mathbf{x})$ as $\sum_{\iota \in C_1} \sum_{k: \iota \in F_k} x_{\iota,k}$. As each inner sum is non-increasing by (7, 8), we conclude that g_1 is non-increasing on all paths. Moreover, for $i \in C_1$, the sum $\sum_{k: i \in F_k} x_{i,k}$ appears in g_1 and is decreasing according to (7), thus g_1 is decreasing for each Q_i with $i \in C_1$.
2. g_1 is non-negative for all Q_i with $i \in C_1$, because all f_k are non-negative on these paths. To see this, pick an arbitrary $i \in C_1$ and an arbitrary face F_k : if $i \in F_k$, and it is the only vertex from C_1 in F_k , then $f_k(\mathbf{x}) = x_{i,k}$ is non-negative on Q_i by (9); if $i \in F_k$ but there is another vertex $j \in C_1$ in F_k , then $f_k(\mathbf{x}) = x_{i,k} + x_{j,k}$ is non-negative on Q_i by (11); if $i \notin F_k$, then for any $j \in F_k$ we have $x_{j,k} \geq 0$ by (10), and then f_k is non-negative since it is a sum of such variables. Note that g_1 can be negative for Q_j with $j \notin C_1$.

Similarly, we construct BMS-QLRFs g_2, \dots, g_d such that g_ℓ ranks Q_i for $i \in C_\ell$. Clearly $\langle g_1, \dots, g_d \rangle$ is a BMS-LLRF for this loop.

Now suppose we have a BMS-LLRF of dimension d ; we analyze what paths Q_i can be associated with each component, and show that for any face F_k , the three paths that are indexed by its vertices, i.e., Q_i for $i \in F_k$, cannot be all associated with the same component. Which clearly yields a d -coloring.

Suppose that for some face $F_k = \{i_1, i_2, i_3\}$, the paths Q_{i_1}, Q_{i_2} and Q_{i_3} are associated with the same component, i.e., all ranked by the same function, say g . Thus $\Delta g(\mathbf{x}'') \geq 1$ must be implied by the constraints of Q_{i_1}, Q_{i_2} and Q_{i_3} , independently. Now since, in each path, the only constraint with a non-zero free coefficient is (7), it follows that the coefficients of variables $x_{i_1,k}, x_{i_2,k}$ and $x_{i_3,k}$ in $g(\mathbf{x})$ are positive, i.e., $g(\mathbf{x}) = a_1 \cdot x_{i_1,k} + a_2 \cdot x_{i_2,k} + a_3 \cdot x_{i_3,k} + h(\mathbf{x})$ where $h(\mathbf{x})$ is a combination of other variables, and $a_1, a_2, a_3 > 0$. Similarly, $g(\mathbf{x}) \geq 0$ must be implied by the constraints of each of three paths independently. For this to hold, g must be a positive linear combination of functions constrained to be non-negative by these paths, and do not involve primed variables. Now consider variables $x_{i_1,k}, x_{i_2,k}$ and $x_{i_3,k}$, and note that they participate only in the following constraints in Q_{i_1} (left), Q_{i_2} (middle) and Q_{i_3} (right):

$$\begin{array}{lll}
 x_{i_1,k} \geq 0 & x_{i_2,k} \geq 0 & x_{i_3,k} \geq 0 \\
 x_{i_1,k} + x_{i_2,k} \geq 0 & x_{i_1,k} + x_{i_2,k} \geq 0 & x_{i_2,k} + x_{i_3,k} \geq 0 \\
 x_{i_1,k} + x_{i_3,k} \geq 0 & x_{i_2,k} + x_{i_3,k} \geq 0 & x_{i_1,k} + x_{i_3,k} \geq 0
 \end{array}$$

This means that the corresponding coefficients in g , i.e., $\bar{a} = (a_1 \ a_2 \ a_3)$, must be equal to linear combinations of the corresponding coefficients in the above

constraints. Namely, there exist $b_1, \dots, b_9 \geq 0$ such that

$$\bar{a} = (b_1 \ b_2 \ b_3) \cdot \begin{pmatrix} 1 & 0 & 0 \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{pmatrix} \quad \bar{a} = (b_4 \ b_5 \ b_6) \cdot \begin{pmatrix} 0 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 1 & 1 \end{pmatrix} \quad \bar{a} = (b_7 \ b_8 \ b_9) \cdot \begin{pmatrix} 0 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 0 & 1 \end{pmatrix}$$

From these nine equations, and the constraints $b_i \geq 0$ for all i , we necessarily get $a_1 = a_2 = a_3 = 0$, which contradicts $a_1, a_2, a_3 > 0$ as we concluded before, and thus paths corresponding to $\{i_1, i_2, i_3\}$ of F_k cannot be all associated with the same component. This concludes the proof of Theorem 3.

6 Complexity of BMS-LEXLINRF (d, \mathbb{Z})

In this section we turn to the problem BMS-LEXLINRF(d, \mathbb{Z}), and show that it is harder than BMS-LEXLINRF(d, \mathbb{Q}), specifically, it is Σ_2^P -complete. The class Σ_2^P is the class of decision problems that can be solved by a standard, non-deterministic computational model in polynomial time assuming access to an oracle for an NP-complete problem. I.e., $\Sigma_2^P = NP^{NP}$. This class contains both NP and coNP, and is likely to differ from them both (this is an open problem).

Theorem 4. *For $d \geq 2$, BMS-LEXLINRF(d, \mathbb{Z}) is a Σ_2^P -complete problem.*

The rest of this section proves Theorem 4. For inclusion in Σ_2^P we use a non-deterministic procedure as in the proof of Theorem 3. Note that the procedure needs to find (or check for existence of) BMS-QLRFs over the integers, so it needs a coNP oracle. For Σ_2^P -hardness we reduce from the canonical Σ_2^P -complete problem (follows from [27, Theorem 4.1]): evaluation of sentences of the form

$$\exists X_1 \dots X_n \ \forall X_{n+1} \dots X_{2n} \ \neg\phi(X_1, \dots, X_{2n}) \tag{*}$$

where the variables X_i are Boolean and the formula ϕ is in 3CNF form. Thus, ϕ is given as a collection of m clauses, C_1, \dots, C_m , each clause C_i consisting of three literals $L_i^j \in \{X_1, \dots, X_{2n}, \neg X_1, \dots, \neg X_{2n}\}$, $1 \leq j \leq 3$. The reduction is first done for $d = 2$, and later extended to $d > 2$ as well.

Let us first explain a well-known approach for reducing satisfiability of a Boolean formula ϕ to satisfiability of integer linear constraints. We first associate each literal L_i^j with an integer variables $x_{i,j}$. Note that the same Boolean variable (or its complement) might be associated with several constraint variables. Let C be the set of (1) all conflicting pairs, that is, pairs $((i, j), (r, s))$ such that L_i^j is the complement of L_r^s ; and (2) pairs $((i, j), (i, j'))$ with $1 \leq j < j' \leq 3$, i.e., pairs of literals that appear in the same clause. We let \mathcal{F} be a conjunction of the constraints: $x_{i,j} + x_{r,s} \leq 1$ for each $((i, j), (r, s)) \in C$; and $0 \leq x_{i,j} \leq 1$ for each $1 \leq i \leq m$ and $1 \leq j \leq 3$. An assignment for $x_{i,j}$ that satisfies \mathcal{F} is called a *non-conflicting assignment*, since if two variables correspond to conflicting literals (or to literals of the same clause) they cannot be assigned 1 at the same time. The next Lemma relates integer assignments with assignments to the Boolean variables of (*). Given a literal L , i.e., X_v or $\neg X_v$, we let $\mathbf{1sum}(L)$ be the sum of all $x_{i,j}$ where $L_i^j \equiv L$ (we use 0 and 1 for *false* and *true*).

Lemma 6. (A) *If σ is a satisfying assignment for ϕ , then there is a non-conflicting assignment for \mathcal{F} such that (1) $x_{i,1} + x_{i,2} + x_{i,3} = 1$ for all $1 \leq i \leq m$; (2) $\sigma(X_v) = 1 \Rightarrow \mathbf{1sum}(\neg X_v) = 0$; and (3) $\sigma(X_v) = 0 \Rightarrow \mathbf{1sum}(X_v) = 0$. (B)* *If ϕ is unsatisfiable, then for any non-conflicting assignment for \mathcal{F} there is at least one $1 \leq i \leq m$ such that $x_{i,1} + x_{i,2} + x_{i,3} = 0$.*

Proof. (A) If σ satisfies ϕ , we construct a satisfying assignment for \mathcal{F} : first every $x_{i,j}$ is assigned the value of L_i^j , and then we turn some $x_{i,j}$ from 1 to 0 so that at most one variable of each clause is set to 1. Since we only turn 1s to 0s, when $\sigma(X_v) = 1$ (resp. $\sigma(X_v) = 0$) all constraint variables that correspond to $\neg X_v$ (resp. X_v) have value 0, and thus $\mathbf{1sum}(\neg X_v) = 0$ (resp. $\mathbf{1sum}(X_v) = 0$). **(B)** If \mathcal{F} has a non-conflicting assignment in which $x_{i,1} + x_{i,2} + x_{i,3} = 1$ for all $1 \leq i \leq m$, then we can construct a satisfying assignment σ for ϕ in which $\sigma(X_v)$ is $\max(\{x_{i,j} | L_i^j \equiv X_v\} \cup \{1 - x_{i,j} | L_i^j \equiv \neg X_v\})$, so ϕ is satisfiable. \square

Next we proceed with the reduction, but first we give an outline. We build an integer loop, call it \mathcal{T} , with $2n + 2$ abstract transitions: $2n$ transitions named $\Psi_{v,a}$, for $1 \leq v \leq n$ and $a \in \{0, 1\}$; plus two named Φ and Ω . These are defined so that existence of a BMS-LLRF $\langle f_1, f_2 \rangle$ for \mathcal{T} implies: (1) $\Psi_{v,0}$ and $\Psi_{v,1}$, for each $1 \leq v \leq n$, cannot be ranked by the same f_i , and the order in which they are ranked will represent a value for the existentially-quantified variable X_v ; (2) Φ cannot be ranked by f_1 , and it is ranked by f_2 iff $\forall X_{n+1} \dots X_{2n} \neg\phi(X_1, \dots, X_{2n})$ is true assuming the values induced for X_1, \dots, X_n in the previous step; and (3) Ω is necessarily ranked by f_1 , its only role is to force Φ to be ranked by f_2 . All these points will imply that (\star) is true. For the other direction, if (\star) is true we show how to construct a BMS-LLRF $\langle f_1, f_2 \rangle$ for \mathcal{T} . Next we formally define the variables and abstract transitions of \mathcal{T} , and prove the above claims.

Variables: Loop \mathcal{T} includes $4m + 2n + 1$ variables: (1) every literal L_i^j contributes a variable $x_{i,j}$; (2) for each $1 \leq i \leq m$, we add a control variable $x_{i,0}$ which is used to check if clause C_i is satisfied; (3) for each $1 \leq v \leq n$, we add variables $z_{v,0}$ and $z_{v,1}$ which help in implementing the existential quantification; and (4) variable w , which helps in ranking the auxiliary transition Ω .

Transitions: First we define Φ , the transition that intuitively checks for satisfiability of $\phi(X_1, \dots, X_{2n})$. It is a conjunction of the following constraints

$$0 \leq x_{i,j} \leq 1 \wedge x'_{i,j} = x_{i,j} \qquad \text{for all } 1 \leq i \leq m, 1 \leq j \leq 3 \quad (12)$$

$$x_{i,j} + x_{r,s} \leq 1 \qquad \text{for all } ((i,j), (r,s)) \in C \quad (13)$$

$$x_{i,0} \geq 0 \wedge x'_{i,0} = x_{i,0} + x_{i,1} + x_{i,2} + x_{i,3} - 1 \qquad \text{for all } 1 \leq i \leq m \quad (14)$$

$$z_{v,0} \geq 0 \wedge z'_{v,0} = z_{v,0} - \mathbf{1sum}(X_v) \qquad \text{for all } 1 \leq v \leq n \quad (15)$$

$$z_{v,1} \geq 0 \wedge z'_{v,1} = z_{v,1} - \mathbf{1sum}(\neg X_v) \qquad \text{for all } 1 \leq v \leq n \quad (16)$$

$$w' = w \qquad (17)$$

Secondly, we define $2n$ transitions which, intuitively, force a choice of a Boolean value for each of X_1, \dots, X_n . For $1 \leq v \leq n$ and $a \in \{0, 1\}$,

transition $\Psi_{v,a}$ is defined as a conjunction of the following constraints

$$z_{v,a} \geq 0 \wedge z'_{v,a} = z_{v,a} - 1 \quad (18)$$

$$z_{u,b} \geq 0 \quad \text{for all } 1 \leq u \leq n, b \in \{0, 1\}, u \neq v \quad (19)$$

$$z'_{u,b} = z_{u,b} \quad \text{for all } 1 \leq u \leq n, b \in \{0, 1\}, (u, b) \neq (v, a) \quad (20)$$

$$x'_{i,0} \geq 0 \wedge x'_{i,0} = x_{i,0} \quad \text{for all } 1 \leq i \leq m \quad (21)$$

$$w \geq 0 \wedge w' = w \quad (22)$$

Finally we define the abstract transition Ω , which aids in forcing a desired form of the BMS-LLRF, and it is defined as a conjunction of the following constraints

$$w \geq 0 \wedge w' = w - 1 \quad (23)$$

$$z_{u,b} \geq 0 \wedge z'_{u,b} = z_{u,b} \quad \text{for all } 1 \leq u \leq n, b \in \{0, 1\} \quad (24)$$

Now, we argue that in order to have a two-component BMS-LLRF for \mathcal{T} , the transitions have to be associated to the two components in a particular way.

Lemma 7. *Suppose that $\langle f_1, f_2 \rangle$ is a BMS-LLRF for \mathcal{T} . Then, necessarily, the correspondence between the BMS-LLRF components and transitions is as follows: (i) Ω is ranked by f_1 ; (ii) Φ is ranked by f_2 ; (iii) for $1 \leq v \leq n$, one of $\Psi_{v,0}$ and $\Psi_{v,1}$ is ranked by f_1 , and the other by f_2 .*

Proof. An LRF for Ω must involve w , since it is the only decreasing variable, and cannot involve any $x_{i,j}$ since they change randomly. Similarly, an LRF for Φ cannot involve w as it has no lower bound, and it must involve at least one $x_{i,j}$ since no function that involves only $z_{v,a}$ variable(s) decreases for an initial state in which all $x_{i,j}$ are assigned 0. Note that such LRF cannot be non-increasing for Ω since $x_{i,j}$ change randomly in Ω . Thus, we conclude that Ω must be associated with f_1 and Φ with f_2 . For the last point, for each $1 \leq v \leq n$, transitions $\Psi_{v,0}$ and $\Psi_{v,1}$ must correspond to different positions because variables that descend in one (namely $z_{v,a}$ of $\Psi_{v,a}$) are not bounded in the other (since (19) requires $u \neq v$). \square

Lemma 8. *A BMS-LLRF of dimension two exists for \mathcal{T} iff (\star) is true.*

Proof. Assume that a BMS-LLRF $\langle f_1, f_2 \rangle$ exists for \mathcal{T} , we show that (\star) is true. By Lemma 7 we know how the transitions are associated with the positions, up to the choice of placing $\Psi_{v,0}$ and $\Psi_{v,1}$, for each $1 \leq v \leq n$. Suppose that, for each $1 \leq v \leq n$, the one which is associated with f_2 is Ψ_{v,a_v} , i.e., $a_v \in \{0, 1\}$, and let \bar{a}_v be the complement of a_v . By construction we know that: (i) in Ψ_{v,a_v} the variables z_{v,\bar{a}_v} and $x_{i,j}$ with $j \geq 1$ change randomly, which means that f_2 cannot involve them; and (ii) in Φ the variable w is not lower bounded, which means that f_2 cannot involve w . Since these transitions must be ranked by f_2 , we can assume that f_2 has the form $f_2(\mathbf{x}, \mathbf{z}, w) = \sum_i c_i \cdot x_{i,0} + \sum_v c_v \cdot z_{v,a_v}$ where c_i and c_v are non-negative rational coefficients. We claim that (\star) is necessarily true; for that purpose we select the value a_v for each X_v , and next we show that this makes it impossible to satisfy $\phi(X_1, \dots, X_{2n})$. Assume, to the contrary, that there is a satisfying assignment σ for ϕ , such that $\sigma(X_v) = a_v$ for all $1 \leq v \leq n$. By Lemma 6 we know that we can construct an assignment to the variables $x_{i,j}$

such that (i) $x_{i,1} + x_{i,2} + x_{i,3} = 1$, for each $1 \leq i \leq m$, which means that $x'_{i,0} = x_{i,0}$ at (14); and (ii) for each $1 \leq v \leq m$, if $a_v = 0$ (resp. $a_v = 1$), then $\mathbf{lsum}(X_v) = 0$ (resp. $\mathbf{lsum}(\neg X_v) = 0$), which means that $z'_{v,a_v} = z_{v,a_v}$ at (15) (resp. (16)). Hence f_2 as described above does *not* rank Φ since none of its variables change, contradicting our assumption. We conclude that (\star) is true.

Now assume that (\star) is true, we construct a BMS-LLRF of dimension two. The assumption means that there are values a_1, \dots, a_n for the existentially-quantified variables to satisfy the sentence. Let $f_1(\mathbf{x}, \mathbf{z}, w) = w + \sum_{v=1}^n z_{v,\bar{a}_v}$ and $f_2(\mathbf{x}, \mathbf{z}, w) = \sum_{i=1}^m x_{i,0} + \sum_v z_{v,a_v}$. We claim that $\langle f_1, f_2 \rangle$ is a BMS-LLRF such that: (i) f_1 is an LRF for Ω and Ψ_{v,\bar{a}_v} , and non-increasing for Ψ_{v,a_v} and Φ ; and (ii) f_2 is an LRF for Ψ_{v,a_v} and Φ . All this is easy to verify, except possibly that f_2 is an LRF for Φ , for which we argue in more detail. By assumption, $\phi(a_1, \dots, a_n, X_{n+1}, \dots, X_{2n})$ is unsatisfiable. Consider a state in which Φ is enabled; by (12, 13), this state may be interpreted as a selection of non-conflicting literals. If one of the selected literals does not agree with the assignment chosen for X_1, \dots, X_n , then by (15, 16) the corresponding variable z_{v,a_v} is decreasing. Otherwise, there must be an unsatisfied clause, and the corresponding variable $x_{i,0}$ is decreasing. All other variables involved in f_2 are non-increasing, all are lower bounded, so f_2 is an LRF for Φ . \square

Σ_2^P -hardness of BMS-LEXLINRF(d, \mathbb{Z}) for $d = 2$ follows from Lemma 8. For $d > 2$, we add to \mathcal{T} additional $d - 2$ paths as those of Example 3; and to each original path in \mathcal{T} we add $x'=x$ and $y'=y$ (x, y are used in Example 3). Then, the new loop has a BMS-LLRF of dimension d iff (\star) is true. This concludes the proof of Theorem 4.

7 Related Work

LLRFs appear in the classic works of Turing [28] and Floyd [16]. Automatic generation of LRFs and LLRFs for linear-constraint loops begins, in the context of logic programs, with Sohn and van Gelder [26]. For imperative programs, it begins with Colón and Sipma [10, 11]. The work of Feautrier on scheduling [14, 15] includes, in essence, generation of LRFs and LLRFs. All these works gave algorithms that yield polynomial time complexity (inherited from LP), except for Colón and Sipma's method which is based on LP duality and polars. The polynomial-time LP method later reappeared in [21, 23]. These methods are complete over the rationals and can be used in an integer setting by relaxing the loop from integer to rational variables, sacrificing completeness. This completeness problem was pointed out (but not solved) in [21, 24], while [12, 14] pointed out the role of the integer hull in ensuring completeness. Bradley et al. [7] use a bisection search over the space of coefficients for inferring LRFs over the integers, which yields completeness at exponential cost (as argued in [4]).

Alias et al. [1] extended the LP approach to LLRFs, obtaining a polynomial-time algorithm which is sound and complete over the rationals (for their notion

of LLRF). The (earlier) work of Bradley et al. [5] introduced BMS-LLRFs and used a “constraint-solving method” that finds such LLRFs along with supporting invariants. The method involves an exponential search for the association of paths to LLRF components, and is complete over the reals. Subsequent work used more complex extensions of the LLRF concept [6, 8]. Harris et al. [17] demonstrate that it is advantageous, to a tool that is based on a CEGAR loop, to search for LLRFs instead of LRFs only. The LLRFs they use are BMS-LLRFs. Similar observations have been reported in [13] (also using BMS-LLRFs), [9] (using ADFG-LLRFs) and [19] (using an iterative construction that extends BMS-LLRFs). Heizmann and Leike [18] generalize the constraint-based approach by defining the concept of a “template” for which one can solve using a constraint solver. They also provide a template for ADFG-LLRFs (of constant dimension). Ben-Amram [3] shows that *every* terminating monotonicity-constraint program has a *piecewise* LLRF of dimension at most $2n$. Piecewise LLRFs are also used in [29], with no completeness result, there they are inferred by abstract interpretation.

8 Conclusion

This work contributes to understanding the design space of the ranking-function method, a well-known method for termination analysis of numeric loops, as well as related analyses (iteration bounds, parallelization schedules). This design space is inhabited by several kinds of “ranking functions” previously proposed. We focused on BMS-LLRFs and compared them to other proposals of a similar nature. We characterized the complexity of finding, or deciding the existence of, BMS-LLRF for rational and integer MLC loops. We also compared these three methods regarding the dimension of the LLRF, and the complexity of optimizing the dimension, which turns out to be essentially harder for BMS-LLRFs. Given our reductions, it is easy to show that it is impossible to approximate the minimal dimension of BMS-LLRFs, in polynomial time, within a factor *smaller than* $\frac{3}{2}$, unless $P=NP$ for rational loops, and $\Sigma_2^P = \Delta_2^P$ for integer loops (see [2]).

We conclude that none of the three methods is universally preferable. Even ADFG-LLRFs, which in principle are weaker than both other methods, have an advantage, in that the algorithm for computing them may be more efficient in practice (due to solving smaller LP problems). If this is not a concern, they can be replaced by BG-LLRFs, so we are left with two, incomparable techniques. This incomparability stems from the fact that BG-LLRFs and BMS-LLRFs relax the restrictions of ADFG-LLRFs in two orthogonal directions: the first in quantifying over concrete transitions rather than abstract ones, and the second in allowing negative components. By making both relaxations, we get a new type of LLRF [19], which is as in Definition 2 but relaxing condition (5) to hold only for $j = i$, but for which the computational complexity questions are still open.

References

1. Alias, C., Darte, A., Feautrier, P., Gonnord, L.: Multi-dimensional rankings, program termination, and complexity bounds of flowchart programs. In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 117–133. Springer, Heidelberg (2010)
2. Ben-Amram, A., Genaim, S.: Complexity of Bradley-Manna-Sipma lexicographic ranking functions. In: CoRR, abs/1504.05018 (2015)
3. Ben-Amram, A.M.: Monotonicity constraints for termination in the integer domain. *Logical Methods in Comput. Sci.* **7**(3) (2011)
4. Ben-Amram, A.M., Genaim, S.: Ranking functions for linear-constraint loops. *J. ACM* **61**(4), 26:1–26:55 (2014)
5. Bradley, A.R., Manna, Z., Sipma, H.B.: Linear ranking with reachability. In: Etessami, K., Rajamani, S.K. (eds.) CAV 2005. LNCS, vol. 3576, pp. 491–504. Springer, Heidelberg (2005)
6. Bradley, A.R., Manna, Z., Sipma, H.B.: The polyranking principle. In: Caires, L., Italiano, G.F., Monteiro, L., Palamidessi, C., Yung, M. (eds.) ICALP 2005. LNCS, vol. 3580, pp. 1349–1361. Springer, Heidelberg (2005)
7. Bradley, A.R., Manna, Z., Sipma, H.B.: Termination Analysis of Integer Linear Loops. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 488–502. Springer, Heidelberg (2005)
8. Bradley, A.R., Manna, Z., Sipma, H.B.: Termination of polynomial programs. In: Cousot, R. (ed.) VMCAI 2005. LNCS, vol. 3385, pp. 113–129. Springer, Heidelberg (2005)
9. Brockschmidt, M., Cook, B., Fuhs, C.: Better termination proving through cooperation. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 413–429. Springer, Heidelberg (2013)
10. Colón, M., Sipma, H.: Synthesis of linear ranking functions. In: Margaria, T., Yi, W. (eds.) TACAS 2001. LNCS, vol. 2031, pp. 67–81. Springer, Heidelberg (2001)
11. Colón, M., Sipma, H.: Practical methods for proving program termination. In: Brinksma, D., Larsen, K.G. (eds.) CAV 2002. LNCS, vol. 2404, pp. 442–454. Springer, Heidelberg (2002)
12. Cook, B., Kroening, D., Rümmer, P., Wintersteiger, C.M.: Ranking function synthesis for bit-vector relations. *Formal Methods in System Design* **43**(1), 93–120 (2013)
13. Cook, B., See, A., Zuleger, F.: Ramsey vs. lexicographic termination proving. In: Piterman, N., Smolka, S.A. (eds.) TACAS 2013 (ETAPS 2013). LNCS, vol. 7795, pp. 47–61. Springer, Heidelberg (2013)
14. Feautrier, P.: Some efficient solutions to the affine scheduling problem. I. One-dimensional time. *Int. J. Parallel Prog.* **21**(5), 313–347 (1992)
15. Feautrier, P.: Some efficient solutions to the affine scheduling problem. II. Multi-dimensional time. *Int. J. Parallel Prog.* **21**(6), 389–420 (1992)
16. Floyd, R.W.: Assigning meanings to programs. *Proc. Symp. Appl. Math.* **XIX**, 19–32 (1967)
17. Harris, W.R., Lal, A., Nori, A.V., Rajamani, S.K.: Alternation for termination. In: Cousot, R., Martel, M. (eds.) SAS 2010. LNCS, vol. 6337, pp. 304–319. Springer, Heidelberg (2010)
18. Leike, J., Heizmann, M.: Ranking templates for linear loops. In: Ábrahám, E., Havelund, K. (eds.) TACAS 2014 (ETAPS). LNCS, vol. 8413, pp. 172–186. Springer, Heidelberg (2014)

19. Larraz, D., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: Proving termination of imperative programs using Max-SMT. In: Formal Methods in Computer-Aided Design, FMCAD 2013, pp. 218–225. IEEE (2013)
20. László, L.: Coverings and colorings of hypergraphs. In: Proceedings of the 4th Southeastern Conference on Combinatorics, Graph Theory, and Computing, pp. 3–12 (1973)
21. Mesnard, F., Serebrenik, A.: Recurrence with affine level mappings is P-time decidable for CLP(R). TPLP **8**(1), 111–119 (2008)
22. Phelps, K.T., Rödl, V.: On the algorithmic complexity of coloring simple hypergraphs and steiner triple systems. *Combinatorica* **4**(1), 79–88 (1984)
23. Podelski, A., Rybalchenko, A.: A complete method for the synthesis of linear ranking functions. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 239–251. Springer, Heidelberg (2004)
24. Rybalchenko, A.: Temporal verification with transition invariants. Ph.D. thesis, Universität des Saarlandes (2004)
25. Schrijver, A.: Theory of Linear and Integer Programming. John Wiley and Sons, New York (1986)
26. Sohn, K., Van Gelder, A.: Termination detection in logic programs using argument sizes. In: Rosenkrantz, D.J. (ed.) Symposium on Principles of Database Systems, pp. 216–226. ACM Press, New York (1991)
27. Stockmeyer, L.J.: The polynomial-time hierarchy. *Theor. Comput. Sci.* **3**(1), 1–22 (1976)
28. Turing, A.M.: Checking a large routine. In: Report of a Conference on High Speed Automatic Calculating Machines, pp. 67–69, 1948. reprinted. In: The early British computer conferences, Charles Babbage Institute Reprint Series For The History Of Computing, vol. 14. MIT Press (1989)
29. Urban, C., Miné, A.: An abstract domain to infer ordinal-valued ranking functions. In: Shao, Z. (ed.) ESOP 2014 (ETAPS). LNCS, vol. 8410, pp. 412–431. Springer, Heidelberg (2014)