

Fast Acceleration of Ultimately Periodic Relations^{*}

Marius Bozga¹, Radu Iosif¹, and Filip Konečný^{1,2}

¹ VERIMAG, CNRS, 2 av. de Vignate, 38610 Gières, France, {bozga, iosif}@imag.fr

² FIT BUT, Božetěchova 2, 61266, Brno, Czech Republic, {ikonecny}@fit.vutbr.cz

Abstract. Computing transitive closures of integer relations is the key to finding precise invariants of integer programs. In this paper, we describe an efficient algorithm for computing the transitive closures of difference bounds, octagonal and finite monoid affine relations. On the theoretical side, this framework provides a common solution to the acceleration problem, for all these three classes of relations. In practice, according to our experiments, the new method performs up to four orders of magnitude better than the previous ones, making it a promising approach for the verification of integer programs.

1 Introduction

The verification of safety properties of infinite-state systems (such as device drivers, communication protocols, control software, etc.) requires the computation of the set of reachable states, starting with an initial state from a given (possibly infinite) set. There are currently two ways of doing this: (i) compute a finite representation of an over-approximation of the set of reachable states, by applying a widening operator at each step, or (ii) attempt to compute precisely the transitive closure of the transition relation; the set of reachable states is the image of the set of initial states via the transitive closure. The first approach is guaranteed to terminate, but the abstraction usually introduces imprecision that may blur the verification result. On the other hand, the second approach, although precise, is not guaranteed to terminate – the problem of verifying safety properties being, in general, undecidable.

In practice, one usually tries to combine the two approaches and benefit from the advantages of both. To this end, it is important to know for which classes of transition relations it is possible to compute the transitive closure precisely and fast – the relations falling outside these classes being dealt with using suitable abstractions. To the best of our knowledge, the three main classes of integer relations for which transitive closures can be computed precisely in finite time are: (1) difference bounds constraints [9, 8], (2) octagons [12, 6], and (3) finite monoid affine transformations [5, 10]. For these three classes, the transitive closures can be moreover defined in Presburger arithmetic.

The contributions of this paper are two-fold. On the theoretical side, we show that the three classes of relations mentioned above are ultimately periodic, i.e. each relation R can be mapped into an integer matrix M_R such that the sequence $\{M_{R^k}\}_{k=0}^{\infty}$ is periodic. The proof that a sequence of matrices is ultimately periodic relies on a result from tropical semiring theory [13]. This provides shorter proofs to the fact that the transitive

^{*} This work was supported by the French national project ANR-09-SEGI-016 VERIDYC, by the Czech Science Foundation (projects P103/10/0306 and 102/09/H042), the Czech Ministry of Education (projects COST OC10009 and MSM 0021630528), and the internal FIT BUT grant FIT-10-1.

closures for these classes can be effectively computed, and that they are Presburger definable.

On the practical side, the algorithm introduced in this paper computes the transitive closure of difference bounds and octagonal relations up to four orders of magnitude faster than the original methods from [8, 6], and also scales much better in the number of variables. The experimental comparison with the FAST tool [4] for difference bounds relations shows that large relations (> 50 variables), causing FAST to run out of memory, can now be handled by our implementation in less than 8 seconds, on average. We currently do not have a full implementation of the finite monoid affine transformation class, which is needed in order to compare our method with tools like FAST [4], LASH [14], or TReX [2], for this class of relations.

Related Work Early attempts to apply Model Checking techniques to the verification of infinite-state systems consider the problem of accelerating transition relations by successive under-approximations, without any guarantee of termination. For systems with integer variables, the acceleration of affine relations has been considered primarily in the works of Annichini et. al [1], Boigelot [5], and Finkel and Leroux [10]. Finite monoid affine relations have been first studied by Boigelot [5], who shows that the finite monoid property is decidable, and that the transitive closure is Presburger definable in this case. On what concerns non-deterministic transition relations, difference bounds constraints appear in the context of timed automata verification. The transitive closure of a difference bounds constraint is shown to be Presburger definable first by Comon and Jurski [9]. Their proof was subsequently simplified and extended to parametric difference bounds constraints in [8]. We also showed that octagonal relations can be accelerated precisely, and that the transitive closure is also Presburger definable [6]. The proofs of ultimate periodicity from this paper are based on some of our previous results [8, 6]. For difference bounds constraints, the proof from [8] was simplified using a result from tropical semiring theory [13].

Roadmap The paper is organized as follows: Section 2 gives the definition of ultimately periodic relations, Section 3 describes the algorithm for computing transitive closures of ultimately periodic relations, in general, Section 4 describes three instances of the algorithm, Section 5 presents the experimental results, and Section 6 concludes. Missing proofs are deferred to [7] due to reasons of space.

2 Preliminaries

We denote by \mathbb{Z} , \mathbb{N} and \mathbb{N}_+ the sets of integers, positive (including zero) and strictly positive integers, respectively. The first order additive theory of integers is known as Presburger Arithmetic. The *tropical semiring* is defined as $\mathbb{T} = (\mathbb{Z}_\infty, \min, +, \infty, 0)$ [13], where $\mathbb{Z}_\infty = \mathbb{Z} \cup \{\infty\}$, with the extended arithmetic operations $x + \infty = \infty$, $\min(x, \infty) = x$, for all $x \in \mathbb{Z}$, where $\min(x, y)$ denotes the minimum between the values x and y . For two square matrices $A, B \in S^{m \times m}$, we define $(A + B)_{ij} = A_{ij} + B_{ij}$ and $(A \times B)_{ij} = \min_{k=1}^m (a_{ik} + b_{kj})$, for all $1 \leq i, j \leq m$. Let $\mathbf{I} \in \mathbb{T}^{m \times m}$ be the identity matrix, i.e. $\mathbf{I}_{ii} = 0$ and $\mathbf{I}_{ij} = \infty$, for all $1 \leq i, j \leq m, i \neq j$.

Definition 1. [13] An infinite sequence $\{s_k\}_{k=0}^\infty \in \mathbb{T}$ is called ultimately periodic if:

$$\exists K \exists c > 0 \exists \lambda_0, \lambda_1, \dots, \lambda_{c-1} \in \mathbb{T} . s_{(k+1)c+i} = \lambda_i + s_{kc+i}$$

for all $k \geq K$ and $i = 0, 1, \dots, c-1$. The smallest c and $\lambda_0, \lambda_1, \dots, \lambda_{c-1}$ for which the above holds are called the period and rates of $\{s_k\}_{k=0}^\infty$, respectively.

Example 1. The sequence $\sigma_k = \{3k+1 \mid k = 2l, l \geq 2\} \cup \{5k+3 \mid k = 2l+1, l \geq 2\}$ is ultimately periodic, with $K = 4$, period $c = 2$ and rates $\lambda_0 = 6, \lambda_1 = 10$. \square

A sequence of matrices $\{A_k\}_{k=0}^\infty \in \mathbb{T}^{m \times m}$ is said to be ultimately periodic if, for all $1 \leq i, j \leq m$, the sequence $\{(A_k)_{ij}\}_{k=0}^\infty$ is ultimately periodic. A matrix $A \in \mathbb{T}^{m \times m}$ is called ultimately periodic if the sequence $\{A^k\}_{k=1}^\infty$ is ultimately periodic, where $A^0 = \mathbf{I}$ and $A^k = A \times A^{k-1}$, for any $k > 0$. It is known that, every matrix is ultimately periodic in the tropical semiring [13].

We have the following characterization of ultimately periodic sequences of matrices:

Lemma 1. *A sequence of matrices $\{A_k\}_{k=1}^\infty \in \mathbb{T}^{m \times m}$ is ultimately periodic if and only if:*

$$\exists K \exists c > 0 \exists \Lambda_0, \Lambda_1, \dots, \Lambda_{c-1} \in \mathbb{T}^{m \times m} . A_{(k+1)c+i} = \Lambda_i + A_{kc+i}$$

for all $k \geq K$ and $i = 0, 1, \dots, c-1$.

If $A \in \mathbb{T}^{m \times m}$ is a square matrix and $n \in \mathbb{T}$, we define the matrix $(n \cdot A)_{ij} = n \cdot A_{ij}$, for all $1 \leq i, j \leq m$. If k is a parameter (typically interpreted over \mathbb{T}), then $\mathbb{T}[k]$ denotes the set of all terms where k may occur, built from the constants and operators of \mathbb{T} . For instance, if $A, B \in \mathbb{T}^{m \times m}$, then $k \cdot A + B \in \mathbb{T}[k]^{m \times m}$ denotes the matrix of terms $(k \cdot A + B)_{ij} = k \cdot A_{ij} + B_{ij}$, for all $1 \leq i, j \leq m$.

2.1 Ultimately Periodic Relations

Let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$ be a set of variables, $N > 0$, and let $\mathbf{x}' = \{x'_1, x'_2, \dots, x'_N\}$. A relation is an arithmetic formula $R(\mathbf{x}, \mathbf{x}')$ with free variables $\mathbf{x} \cup \mathbf{x}'$. We say that two relations R and R' are equivalent, denoted $R \Leftrightarrow R'$ if under all valuations of \mathbf{x} and \mathbf{x}' , R is true if and only if R' is true. A relation is called *consistent* if and only if there exist valuations of \mathbf{x} and \mathbf{x}' under which it holds. We denote a consistent relation R by writing $R \not\Leftarrow$ false, and an inconsistent relation by writing $R \Leftarrow$ false.

The composition of two relations is defined as $R \circ R' \equiv \exists \mathbf{y} . R(\mathbf{x}, \mathbf{y}) \wedge R'(\mathbf{y}, \mathbf{x}')$. Let \mathcal{I} be the identity relation $\bigwedge_{x \in \mathbf{x}} x' = x$. We define $R^0 \equiv \mathcal{I}$ and $R^n \equiv R^{n-1} \circ R$, for any $n > 0$. With these notations, $R^* \equiv \bigvee_{i=0}^\infty R^i$ denotes the *transitive closure* of R . A relation R is called ω -consistent if R^n is consistent for all $n > 0$. For the rest of this section, let \mathcal{R} be a class of relations³.

Definition 2. *A relation $R(\mathbf{x}, \mathbf{x}') \in \mathcal{R}$ is called ultimately periodic if and only if either:*

1. *there exists $i_0 \geq 0$ such that R^{i_0} is inconsistent, or*
2. *for all $i \geq 0$, R^i is consistent, and there exists two functions:*
 - $\sigma : \mathcal{R} \rightarrow \mathbb{T}_{\perp}^{m \times m}$ *mapping each consistent relation in \mathcal{R} into a $m \times m$ matrix of \mathbb{T} , for some $m > 0$, and each inconsistent relation into \perp .*
 - $\rho : \mathbb{T}^{m \times m} \rightarrow \mathcal{R}$ *mapping each $m \times m$ matrix of \mathbb{T} into a relation in \mathcal{R} , such that $\rho(\sigma(R)) \Leftrightarrow R$, for each consistent relation $R \in \mathcal{R}$*

³ A class of relations is usually defined by syntactic conditions.

such that the infinite sequence of matrices $\{\sigma(R^i)\}_{i=0}^{\infty} \in \mathbb{T}^{m \times m}$ is ultimately periodic.

Notice that the first condition of the definition implies that $\sigma(R^i) = \perp$, for all $i \geq i_0$. If each relation $R \in \mathcal{R}$ is ultimately periodic, then \mathcal{R} is called ultimately periodic as well. The following lemma gives an alternative characterization of ω -consistent ultimately periodic relations.

Lemma 2. *An ω -consistent relation R is ultimately periodic if and only if there exist $K \geq 0$, $b \geq 0$, $c > 0$ and $\Lambda_0, \Lambda_1, \dots, \Lambda_{c-1} \in \mathbb{T}^{m \times m}$ such that the following hold:*

1. $\sigma(R^{(n+1)c+i}) = \Lambda_i + \sigma(R^{nc+i})$, for all $n \geq K$.
2. $R^{nc+b+i} \Leftrightarrow \rho(n \cdot \Lambda_i + \sigma(R^{b+i}))$, for all $n \geq 0$.

for all $i = 0, 1, \dots, c-1$, where σ and ρ are the functions from Def. 2.

Proof. By Lemma 1, if R is ω -consistent, then it is ultimately periodic if and only if

$$\exists K \exists c > 0 \exists \Lambda_0, \Lambda_1, \dots, \Lambda_{c-1} \in \mathbb{T}^{N \times N} . \sigma(R^{(k+1)c+i}) = \Lambda_i + \sigma(R^{kc+i})$$

for all $k \geq K$ and $i = 0, 1, \dots, c-1$. By induction on $k \geq K$, one shows first that

$$R^{kc+i} \Leftrightarrow \rho(\Lambda_i^{k-K} + \sigma(R^{Kc+i})), \forall k \geq K$$

Let $b = Kc$. By replacing $k - K$ with k , we obtain

$$R^{kc+b+i} \Leftrightarrow \rho(\Lambda_i^k + \sigma(R^{b+i})), \forall k \geq 0$$

□

For practical reasons related to the representation of R^* , we are interested in finding the symbolic expression R^k , where k is a parameter (because $R^* \equiv \exists k . R^k$). Notice that the second point of lemma 2 can be used to compute the expression R^k symbolically (as a formula over \mathbf{x} , \mathbf{x}' and k), assuming that we are given a function, call it $\pi : \mathbb{T}[k]^{m \times m} \rightarrow \mathcal{R}(k)$, where $\mathcal{R}(k)$ is the class of all parametric relations over \mathbf{x} , \mathbf{x}' and k . Intuitively, π is the parametric counterpart of the ρ function from Def. 2, mapping a matrix of terms over k into a parametric relation $R(\mathbf{x}, \mathbf{x}', k)$. Concrete definitions of π will be given in Section 4.

3 Computing Transitive Closures of Ultimately Periodic Relations

In this section we give a generic algorithm that computes the transitive closure of a given ultimately periodic relation. The algorithm needs to be instantiated for a specific class \mathcal{R} of ultimately periodic relations by providing the mappings σ , ρ (Def. 2) and π (the parametric counterpart of ρ) as discussed in the previous. Next, in Section 4, we show how this algorithm can be used for accelerating three classes of relations: difference bounds, octagons, and finite monoid affine transformations.

Fig. 1 shows the generic framework for computing transitive closures. The input to the algorithm is a relation R , and the mappings $\sigma : \mathcal{R} \rightarrow \mathbb{T}^{m \times m}$, $\rho : \mathbb{T}^{m \times m} \rightarrow \mathcal{R}$, and $\pi : \mathbb{T}[k]^{m \times m} \rightarrow \mathcal{R}(k)$. The algorithm is guaranteed to terminate if R is ultimately periodic, as it will be explained in the following.

The main idea of the algorithm is to discover the prefix b and period c of the sequence $\{\sigma(R^i)\}_{i=0}^{\infty}$ – cf. the second point of lemma 2. If R is ultimately periodic, such values are guaranteed to exist. The dove-tailing enumeration on lines 1 and 2 is guaranteed to yield the smallest value c for which the sequence is shown to be periodic⁴.

Second, the algorithm attempts to compute the first rate of the sequence (line 6), by comparing the matrices $\sigma(R^b)$, $\sigma(R^{c+b})$ and $\sigma(R^{2c+b})$. If the difference Λ between $\sigma(R^{c+b})$ and $\sigma(R^b)$ equals the difference between $\sigma(R^{2c+b})$ and $\sigma(R^{c+b})$, then Λ is a valid candidate for the first rate of the progression (see lemma 2). Notice that the consistency check on line 4 is needed to ensure that we apply σ to consistent relations – otherwise, the relation is not ω -consistent, and the algorithm returns directly the transitive closure, i.e. the finite disjunction $\bigvee_{i=0}^{kc+b-1} R^i$, $0 \leq k \leq 2$ (line 4).

Once a candidate Λ for the initial rate was found, the test \mathcal{Q}_1 on line 7 is used to check that R is ultimately periodic and ω -consistent. Notice that the characterization of ultimately periodic relations from lemma 2 cannot be applied here, since R^n is not known in general, for arbitrary $n \geq 0$. The condition used here is local, i.e. it needs only the relation R^b , for a typically small constant $b \geq 0$. The next lemma establishes the correctness of the criterion used by \mathcal{Q}_1 :

Lemma 3. *An ω -consistent relation R is ultimately periodic if and only if*

$$\exists b \exists c > 0 \exists \Lambda_0, \Lambda_1, \dots, \Lambda_{c-1} \in \mathbb{T}^{m \times m} . \rho(n \cdot \Lambda_i + \sigma(R^{b+i})) \circ R^c \Leftrightarrow \rho((n+1) \cdot \Lambda_i + \sigma(R^{b+i}))$$

for all $n \geq 0$ and $i = 0, 1, \dots, c-1$, where σ and ρ are the functions from Def. 2. Moreover, $\Lambda_0, \Lambda_1, \dots, \Lambda_{c-1}$ satisfy the equivalences of Lemma 2.

Proof. “ \Rightarrow ” If R is ω -consistent and ultimately periodic, by Lemma 2, there exist $b \geq 0$, $c > 0$ and $\Lambda_0, \Lambda_1, \dots, \Lambda_{c-1} \in \mathbb{T}^{m \times m}$ such that

$$R^{kc+b+i} \Leftrightarrow \rho(\Lambda_i^k + \sigma(R^{b+i}))$$

for all $k \geq 0$ and $i = 0, 1, \dots, c-1$. We have:

$$\begin{aligned} R^{(k+1)c+b+i} &\Leftrightarrow R^{kc+b+i} \circ R^c \\ \rho(\Lambda_i^{k+1} + \sigma(R^{b+i})) &\Leftrightarrow \rho(\Lambda_i^k + \sigma(R^{b+i})) \circ R^c \end{aligned}$$

“ \Leftarrow ” We prove the equivalent condition of Lemma 2 by induction on $k \geq 0$. The base case $k = 0$ is immediate. The induction step is as follows:

$$\begin{aligned} R^{(k+1)c+b+i} &\Leftrightarrow R^{kc+b+i} \circ R^c \\ &\Leftrightarrow \rho(\Lambda_i^k + \sigma(R^{b+i})) \circ R^c, \text{ by the induction hypothesis} \\ &\Leftrightarrow \rho(\Lambda_i^{k+1} + \sigma(R^{b+i})) \end{aligned}$$

□

The universal query \mathcal{Q}_1 on line 7 is in general handled by procedures that are specific to the class of relations \mathcal{R} we work with. Notice furthermore that \mathcal{Q}_1 can be handled symbolically by checking the validity of the first order formula: $\forall k . \pi(k \cdot \Lambda + \sigma(R^b)) \circ R^c \Leftrightarrow \pi((k+1) \cdot \Lambda + \sigma(R^b))$, where π is the parametric counterpart of ρ . Next, in Section 4, we detail two ways in which this test can be performed efficiently (for difference bounds and octagonal relations), without resorting to external proof engines, such as SMT or Presburger solvers.

⁴ The nested loop from Fig. 1 will always yield a pair (b, c) such that $b \geq c$. To ensure that b is also minimal, and thus cover up the case $b < c$, once the smallest period c has been detected at prefix $b = c$, we need to also try all prefixes $b = c-1, c-2, \dots, 0$.

```

1. foreach  $b := 0, 1, 2, \dots$  do
2.   foreach  $c := 1, 2, \dots, b$  do
3.     foreach  $k := 0, 1, 2$  do
4.       if  $R^{kc+b} \Leftrightarrow \text{false}$  then return  $R^* \equiv \bigvee_{i=0}^{kc+b-1} R^i$ 
5.     endfor
6.     if exists  $\Lambda \in \mathbb{T}^{m \times m} : \sigma(R^{c+b}) = \Lambda + \sigma(R^b)$  and  $\sigma(R^{2c+b}) = \Lambda + \sigma(R^{c+b})$  then
7.       if forall  $n \geq 0 : \rho(n \cdot \Lambda + \sigma(R^b)) \circ R^c \Leftrightarrow \rho((n+1) \cdot \Lambda + \sigma(R^b)) \Leftrightarrow \text{false}$  ( $\mathcal{Q}_1$ ) then
8.         return  $R^* \equiv \bigvee_{i=0}^{b-1} R^i \vee \exists k \geq 0 . \bigvee_{i=0}^{c-1} \pi(k \cdot \Lambda + \sigma(R^b)) \circ R^i$ 
9.       else if exists  $n \geq 0 : \rho(n \cdot \Lambda + \sigma(R^b)) \circ R^c \Leftrightarrow \text{false}$  ( $\mathcal{Q}_2$ ) then
10.        let  $n_0 = \min\{n \mid \rho(n \cdot \Lambda + \sigma(R^b)) \circ R^c \Leftrightarrow \text{false}\}$ 
11.        if forall  $n \in [0, n_0 - 1] : \rho(n \cdot \Lambda + \sigma(R^b)) \circ R^c \Leftrightarrow \rho((n+1) \cdot \Lambda + \sigma(R^b))$  then
12.          return  $R^* \equiv \bigvee_{i=0}^{b-1} R^i \vee \bigvee_{n=0}^{n_0-1} \bigvee_{i=0}^{c-1} \rho(n \cdot \Lambda + \sigma(R^b)) \circ R^i$ 
13.        endif
14.      endif
15.    endfor
16.  endfor

```

Fig. 1: Transitive Closure Algorithm

If the universal query on line 7 holds, the rate Λ can be used now to express the transitive closure (line 8) as a finite disjunction over the prefix $(\bigvee_{i=0}^{b-1} R^i)$ followed by a formula defining an arbitrary number of iterations $(\exists k . \bigvee_{i=0}^{c-1} \pi(k \cdot \Lambda + \sigma(R^b)) \circ R^i)$. Note that the formula on line 8 defines indeed the transitive closure of R , as a consequence of lemma 2. Moreover, this is a formula of Presburger arithmetic, provided that the classes of relations \mathcal{R} and $\mathcal{R}(k)$ are Presburger definable.

Otherwise, if \mathcal{Q}_1 does not hold, there are two possibilities: either (i) Λ is actually not the first rate of the sequence $\{\sigma(R^i)\}_{i=0}^{\infty}$ for given $b \geq 0$ and $c > 0$, or (ii) the relation is not ω -consistent. In the first case, we need to reiterate with another prefix-period pair, which will give us another candidate Λ .

In the second case, R^m becomes inconsistent, for some $m > 0$ – in this case the computation of its transitive closure is possible, in principle, by taking the disjunction of all powers of R up to m . However, in practice this may take a long time, if m is large. In order to speed up the computation, we check whether:

- $\rho(n \cdot \Lambda + \sigma(R^b)) \circ R^c$ is inconsistent (line 9); the existential query \mathcal{Q}_2 (and namely finding the smallest value for which it holds) is dealt with in Section 4, specifically for the classes of difference bounds and octagonal relations.
- R is periodic with first rate Λ between 0 and $n_0 - 1$ (line 11), where n_0 is the smallest n satisfying the first point (line 10).

If both conditions above hold, then $m = (n_0 + 1)c + b$ is the smallest value for which R^m becomes inconsistent, and moreover, R is periodic with rate Λ between 0 and m . If this is the case, we compute the transitive closure using the period Λ and return (line 12). The following theorem can be proved along the lines of the discussion above:

Theorem 1. *If R is an ultimately periodic relation, the algorithm in Fig. 1 eventually terminates and returns the transitive closure of R .*

4 Some Ultimately Periodic Classes of Arithmetic Relations

This section is dedicated to the application of the transitive closure computation algorithm from the previous section (Fig. 1) to three classes of arithmetic relations, for which the transitive closure is Presburger-definable: difference bounds relations [8], octagonal relations [6], and finite monoid affine transformations [5].

In order to apply the transitive closure computation method, one needs to address two issues. First, the class of relations considered needs to be proved ultimately periodic (or else, our algorithm is not guaranteed to terminate). The proofs rely mostly on the fact that any matrix A is ultimately periodic in \mathbb{T} [13] (see Section 2 for the definition of ultimately periodic matrices).

Second, the queries \mathcal{Q}_1 and \mathcal{Q}_2 (Fig. 1) need to be answered efficiently, by avoiding excessive calls to external decision procedures. In theory, all these queries can be expressed in Presburger arithmetic, for the classes of difference constraints, octagons and affine transformations, yet in practice we would like to avoid as much as possible using Presburger solvers, due to reasons of high complexity. For the classes of difference bounds and octagons, we give direct decision methods for handling these queries. The class of affine transformations without guards can also be dealt with by simply checking equivalence between Diophantine systems, whereas the general case still needs to be handled by a Presburger solver.

4.1 Difference Constraints

Let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$ be a set of variables ranging over \mathbb{Z} .

Definition 3. A formula $\phi(\mathbf{x})$ is a difference bounds constraint if it is equivalent to a finite conjunction of atomic propositions of the form $x_i - x_j \leq a_{ij}$, $1 \leq i, j \leq N, i \neq j$, where $a_{ij} \in \mathbb{Z}$.

For example, $x = y + 5$ is a difference bounds constraint, as it is equivalent to $x - y \leq 5 \wedge y - x \leq -5$. Let \mathcal{R}_{db} denote the class of difference bound relations. Difference bounds constraints are alternatively represented as matrices or, equivalently, weighted graphs.

Given a difference bounds constraint ϕ , a *difference bounds matrix* (DBM) representing ϕ is a matrix $M_\phi \in \mathbb{T}^{N \times N}$ such that $(M_\phi)_{ij} = a_{ij}$, if $x_i - x_j \leq a_{ij}$ is an atomic proposition in ϕ , and ∞ , otherwise. Dually, if $M \in \mathbb{T}^{N \times N}$ is a DBM, the corresponding difference bounds constraint is $\Delta_M \equiv \bigwedge_{M_{ij} < \infty} x_i - x_j \leq M_{ij}$.

A DBM M is said to be consistent if and only if its corresponding constraint φ_M is consistent. An *elementary path* in a DBM M is a sequence of indices $1 \leq i_1, i_2, \dots, i_k \leq N$, where i_1, \dots, i_{k-1} are pairwise distinct, such that $M_{i_j i_{j+1}} < \infty$, for all $1 \leq j < k$. An elementary path is called an *elementary cycle* if moreover $i_1 = i_k$. An elementary cycle is said to be *strictly negative* if $\sum_{j=1}^{k-1} M_{i_j i_{j+1}} < 0$. A DBM M is inconsistent if and only if it has a strictly negative elementary cycle – a proof can be found in [12]. The next definition gives a canonical form for consistent DBMs.

Definition 4. A consistent DBM $M \in \mathbb{T}^{N \times N}$ is said to be closed if and only if $M_{ii} = 0$ and $M_{ij} \leq M_{ik} + M_{kj}$, for all $1 \leq i, j, k \leq N$.

Given a consistent DBM M , we denote by M^* the (unique) closed DBM such that $\varphi_M \Leftrightarrow \varphi_{M^*}$. It is well-known that, if M is consistent, then M^* is unique, and can be computed from M in time $\mathcal{O}(N^3)$, by the classical Floyd-Warshall algorithm. Moreover, if M is a consistent DBM, we have, for all $1 \leq i, j \leq N$:

$$M_{ij}^* = \min \left\{ \sum_{l=0}^{k-1} M_{i_l i_{l+1}} \mid i = i_0 \dots i_{k-1} = j \text{ is an elementary path in } M \right\} \quad (1)$$

The closed form of DBMs is needed for the elimination of existentially quantified variables – if ϕ is a difference bounds constraint, then $\exists x . \phi$ is also a difference bounds constraint [12]. Consequently, we have that the class of difference bounds relations is closed under relational composition: $R_1(\mathbf{x}, \mathbf{x}') \circ R_2(\mathbf{x}, \mathbf{x}') \equiv \exists \mathbf{y} . R_1(\mathbf{x}, \mathbf{y}) \wedge R_2(\mathbf{y}, \mathbf{x}')$.

Difference Bounds Relations are Ultimately Periodic Given a consistent difference bounds relation $R(\mathbf{x}, \mathbf{x}') \in \mathcal{R}_{db}$, let $\sigma(R) = M_R \in \mathbb{T}^{2N \times 2N}$ be the characteristic DBM of R , and for any $M \in \mathbb{T}^{2N \times 2N}$, let $\rho(M) = \Delta_M \in \mathcal{R}_{db}$ be the difference bounds relation corresponding to R . Clearly, $\rho(\sigma(R)) \Leftrightarrow R$, as required by Def. 2.

With these definitions, the algorithm in Fig. 1 will return the transitive closure of a difference bounds relation R , provided that the sequence $\{\sigma(R^i)\}_{i=0}^{\infty}$ is ultimately periodic. If R is not ω -consistent then, by Def. 2, it is ultimately periodic. We consider henceforth that R is ω -consistent, i.e. $\sigma(R^i) = M_{R^i}$, for all $i \geq 0$.

For a difference bounds relation R , we define the directed graph \mathcal{G}_R , whose set of vertices is the set $\mathbf{x} \cup \mathbf{x}'$, and in which there is an edge from x_i to x_j labeled α_{ij} if and only if the atomic proposition $x_i - x_j \leq \alpha_{ij}$ occurs in R . Clearly, M_R is the incidence matrix of \mathcal{G}_R .

Next, we define the concatenation of \mathcal{G}_R with itself as the disjoint union of two copies of \mathcal{G}_R , in which the \mathbf{x}' vertices of the second copy overlap with the \mathbf{x} vertices of the first copy. Then R^m corresponds to the graph \mathcal{G}_R^m , obtained by concatenating the graph of R to itself $m > 0$ times. Since \mathcal{R}_{db} is closed under relational composition, then $R^m \in \mathcal{R}_{db}$, and moreover we have:

$$\bigwedge_{1 \leq i, j \leq N} x_i - x_j \leq \min\{x_i^0 \rightarrow x_j^0\} \wedge x'_i - x'_j \leq \min\{x_i^m \rightarrow x_j^m\} \wedge \\ x_i - x'_j \leq \min\{x_i^0 \rightarrow x_j^m\} \wedge x'_i - x_j \leq \min\{x_i^m \rightarrow x_j^0\}$$

where $\min\{x_i^p \rightarrow x_j^q\}$ is the minimal weight of all paths between the extremal vertices x_i^p and x_j^q in \mathcal{G}_R^m , for $p, q \in \{0, m\}$. In other words, we have the equalities from Fig. 2 (a), for all $1 \leq i, j \leq N$.

$$\begin{array}{ll} (M_{R^m})_{i,j} &= \min\{x_i^0 \rightarrow x_j^0\} & \min\{x_i^0 \rightarrow x_j^0\} &= (\mathcal{M}_R^m)_{I_{ef}(x_i), F_{ef}(x_j)} \\ (M_{R^m})_{i+N, j+N} &= \min\{x_i^m \rightarrow x_j^m\} & \min\{x_i^m \rightarrow x_j^m\} &= (\mathcal{M}_R^m)_{I_{eb}(x_i), F_{eb}(x_j)} \\ (M_{R^m})_{i, j+N} &= \min\{x_i^0 \rightarrow x_j^m\} & \min\{x_i^0 \rightarrow x_j^m\} &= (\mathcal{M}_R^m)_{I_{of}(x_i), F_{of}(x_j)} \\ (M_{R^m})_{i+N, j} &= \min\{x_i^m \rightarrow x_j^0\} & \min\{x_i^m \rightarrow x_j^0\} &= (\mathcal{M}_R^m)_{I_{ob}(x_i), F_{ob}(x_j)} \end{array} \quad \begin{array}{l} (a) \\ (b) \end{array}$$

Fig. 2

As proved in [8], the paths between x_i^p and x_j^q , for arbitrary $1 \leq i, j \leq N$ and $p, q \in \{0, m\}$, can be seen as words (over a finite alphabet of subgraphs of \mathcal{G}_R^m) recognized by a finite weighted automaton of size up to 5^N . For space reasons, the definition of this automaton is detailed in [7].

Let \mathcal{M}_R be the incidence matrix of this automaton. By the construction of \mathcal{M}_R , for each variable $x \in \mathbf{x}$, there are eight indices, denoted as⁵ $I_{of}(x), I_{ob}(x), I_{ef}(x), I_{eb}(x), F_{of}(x), F_{ob}(x), F_{ef}(x), F_{eb}(x) \in \{1, \dots, 5^N\}$, such that all relations from Fig. 2 (b)

⁵ Paths between x^0 and y^m (x^m and y^0) are called odd forward (backward) in [8], whereas paths between x^0 and y^0 (x^m and y^m) are called even forward (backward). Hence the indices *of*, *ob*, *ef* and *eb*.

hold, for all $1 \leq i, j \leq N$. Intuitively, all paths from x_i^0 to x_j^0 are recognized by the automaton with $I_{ef}(x_i)$ and $F_{ef}(x_j)$ as the initial and final states, respectively. The same holds for the other pairs of indices, from Fig. 2 (b). It is easy to see (as an immediate consequence of the interpretation of the matrix product in \mathbb{T}) that, for any $m > 0$, the matrix \mathcal{M}_R^m gives the minimal weight among all paths, of length m , between x_i^p and x_j^q , for any $1 \leq i, j \leq N$ and $p, q \in \{0, m\}$. But the sequence $\{\mathcal{M}_R^m\}_{m=0}^\infty$ is ultimately periodic, since every matrix is ultimately periodic in \mathbb{T} [13]. By equating the relations from Fig. 2 (a) with the ones from Fig. 2 (b), we obtain that the sequence $\{\sigma(R^m)\}_{m=0}^\infty = \{M_{R^m}\}_{m=0}^\infty$ is ultimately periodic as well.

In conclusion, the algorithm from Fig. 1 will terminate on difference bounds relations. Moreover, the result is a formula in Presburger arithmetic. This also simplifies the proof that transitive closures of difference bounds relations are Presburger definable, from [8], since the minimal paths of length m within the weighted automaton recognizing the paths of \mathcal{G}_R^m correspond in fact to elements of the m -th power of \mathcal{M}_R (the incidence matrix of the automaton) in \mathbb{T} .

Checking ω -Consistency and Inconsistency of Difference Bounds Relations For a difference bounds relation $R(\mathbf{x}, \mathbf{x}') \in \mathcal{R}_{db}$ and a matrix $\Lambda \in \mathbb{T}^{2N \times 2N}$, we give methods to decide the queries \mathcal{Q}_1 and \mathcal{Q}_2 (lines 7 and 9 in Fig. 1) efficiently. To this end, we consider the class of parametric difference bounds relations. From now on, let $k \notin \mathbf{x}$ be a variable interpreted over \mathbb{N}_+ .

Definition 5. A formula $\phi(\mathbf{x}, k)$ is a parametric difference bounds constraint if it is equivalent to a finite conjunction of atomic propositions of the form $x_i - x_j \leq a_{ij} \cdot k + b_{ij}$, for some $1 \leq i, j \leq N$, $i \neq j$, where $a_{ij}, b_{ij} \in \mathbb{Z}$.

The class of parametric difference bounds relations with parameter k is denoted as $\mathcal{R}_{db}(k)$. A parametric difference bounds constraint $\phi(k)$ can be represented by a matrix $M_\phi[k]$ of linear terms, where $(M_\phi[k])_{ij} = a_{ij} \cdot k + b_{ij}$ if $x_i - x_j \leq a_{ij} \cdot k + b_{ij}$ occurs in ϕ , and ∞ otherwise. Dually, a matrix $M[k]$ of linear terms corresponds to the formula $\Delta_M(k) \equiv \bigwedge_{M[k]_{ij} \neq \infty} x_i - x_j \leq M[k]_{ij}$. With these considerations, we define $\pi(M[k]) = \Delta_M(k)$. Clearly, we have $\pi(k \cdot \Lambda + \sigma(R^b)) \in \mathcal{R}_{db}(k)$, for $R \in \mathcal{R}_{db}$, $b \geq 0$ and $\Lambda \in \mathbb{T}^{2N \times 2N}$.

Parametric DBMs do not have a closed form, since in general, the minimum of two linear terms in k (for all valuations of k) cannot be expressed again as a linear term. According to (1), one can define the closed form of a parametric DBM as a matrix of terms of the form $\min\{a_i \cdot k + b_i\}_{i=1}^m$, for some $a_i, b_i \in \mathbb{Z}$ and $m > 0$. Then the query \mathcal{Q}_1 can be written as a conjunction of formulae of the form $\forall k > 0. \min\{a_i \cdot k + b_i\}_{i=1}^m = a_0 \cdot k + b_0$. The following lemma gives a way to decide the validity of such formulae:

Lemma 4. Given $\ell, a_0, a_1, \dots, a_m, b_0, b_1, \dots, b_m \in \mathbb{Z}$, the following are equivalent:

1. $\forall k \geq \ell. \min\{a_i \cdot k + b_i\}_{i=1}^m = a_0 \cdot k + b_0$
2. $\bigvee_{i=1}^m (a_i = a_0 \wedge b_i = b_0) \wedge \bigwedge_{j=1}^m (a_0 \leq a_j \wedge a_0 \cdot \ell + b_0 \leq a_j \cdot \ell + b_j)$

In analogy to the non-parametric case, the inconsistency of a parametric difference bounds constraint $\phi(k)$ amounts to the existence of a strictly negative elementary cycle in $M_\phi[k]$, for some valuation $k \in \mathbb{N}_+$. We are also interested in finding the smallest value for which such a cycle exists. The following lemma gives this value.

Lemma 5. Let $\phi(\mathbf{x}, k)$ be a parametric difference bounds constraint and $M_\phi[k]$ be its associated matrix. For some $a_{ij}, b_{ij} \in \mathbb{Z}$, let $\{a_{ij} \cdot k + b_{ij}\}_{j=1}^{m_i}$, $i = 1, \dots, 2N$ be the set of terms denoting weights of elementary cycles going through i . Then ϕ is inconsistent for some $\ell \in \mathbb{N}$ and $k \geq \ell$ if and only if there exists $1 \leq i \leq 2N$ and $1 \leq j \leq m_i$ such that either (i) $a_{ij} < 0$ or (ii) $a_{ij} \geq 0 \wedge a_{ij} \cdot \ell + b_{ij} < 0$ holds. Moreover, the smallest value for which ϕ becomes inconsistent is $\min_{i=1}^{2N} \{\min_{j=1}^{m_i} \gamma_{ij}\}$, where $\gamma_{ij} = \max(\ell, \lfloor -\frac{b_{ij}}{a_{ij}} \rfloor + 1)$, if $a_{ij} < 0$, $\gamma_{ij} = \ell$, if $a_{ij} \geq 0 \wedge a_{ij} \cdot \ell + b_{ij} < 0$, and $\gamma_{ij} = \infty$, otherwise.

4.2 Octagons

Let $\mathbf{x} = \{x_1, x_2, \dots, x_N\}$ be a set of variables ranging over \mathbb{Z} .

Definition 6. A formula $\phi(\mathbf{x})$ is an octagonal constraint if it is equivalent to a finite conjunction of terms of the form $\pm x_i \pm x_j \leq a_{ij}$, $2x_i \leq b_i$, or $-2x_i \leq c_i$, where $a_{ij}, b_i, c_i \in \mathbb{Z}$ and $1 \leq i, j \leq N$, $i \neq j$.

The class of octagonal relations is denoted by \mathcal{R}_{oct} in the following. We represent octagons as difference bounds constraints over the set of variables $\mathbf{y} = \{y_1, y_2, \dots, y_{2N}\}$, with the convention that y_{2i-1} stands for x_i and y_{2i} for $-x_i$, respectively. For example, the octagonal constraint $x_1 + x_2 = 3$ is represented as $y_1 - y_4 \leq 3 \wedge y_2 - y_3 \leq -3$. To handle the \mathbf{y} variables in the following, we define $\bar{i} = i - 1$, if i is even, and $\bar{i} = i + 1$ if i is odd. Obviously, we have $\bar{\bar{i}} = i$, for all $i \in \mathbb{Z}$, $i \geq 0$. We denote by $\bar{\phi}$ the difference bounds formula $\phi[y_1/x_1, y_2/-x_1, \dots, y_{2n-1}/x_n, y_{2n}/-x_n]$ over \mathbf{y} . The following equivalence relates ϕ and $\bar{\phi}$:

$$\phi(\mathbf{x}) \Leftrightarrow (\exists y_2, y_4, \dots, y_{2N} \cdot \bar{\phi} \wedge \bigwedge_{i=1}^N y_{2i-1} + y_{2i} = 0)[x_1/y_1, \dots, x_n/y_{2N-1}] \quad (2)$$

An octagonal constraint ϕ is equivalently represented by the DBM $M_{\bar{\phi}} \in \mathbb{T}^{2N \times 2N}$, corresponding to $\bar{\phi}$. We say that a DBM $M \in \mathbb{T}^{2N \times 2N}$ is *coherent* iff $M_{ij} = M_{j\bar{i}}$ for all $1 \leq i, j \leq 2N$. This property is needed since any atomic proposition $x_i - x_j \leq a$, in ϕ can be represented as both $y_{2i-1} - y_{2j-1} \leq a$ and $y_{2j} - y_{2i} \leq a$, $1 \leq i, j \leq N$. Dually, a coherent DBM $M \in \mathbb{T}^{2N \times 2N}$ corresponds to the octagonal constraint Ω_M :

$$\bigwedge_{1 \leq i, j \leq N} (x_i - x_j \leq M_{2i-1, 2j-1} \wedge x_i + x_j \leq M_{2i-1, 2j} \wedge -x_i - x_j \leq M_{2i, 2j-1}) \quad (3)$$

A coherent DBM M is said to be *octagonal-consistent* if and only if Ω_M is consistent.

Definition 7. An octagonal-consistent coherent DBM $M \in \mathbb{T}^{2N \times 2N}$ is said to be *tightly closed* if and only if the following hold:

1. $M_{ii} = 0$, $\forall 1 \leq i \leq 2N$
2. $M_{i\bar{i}}$ is even, $\forall 1 \leq i \leq 2N$
3. $M_{ij} \leq M_{ik} + M_{kj}$, $\forall 1 \leq i, j, k \leq 2N$
4. $M_{ij} \leq \lfloor \frac{M_{i\bar{i}}}{2} \rfloor + \lfloor \frac{M_{j\bar{j}}}{2} \rfloor$, $\forall 1 \leq i, j \leq 2N$

The following theorem from [3] provides an effective way of testing consistency and computing the tight closure of a coherent DBM. Moreover, it shows that the tight closure of a given DBM is unique and can also be computed in time $\mathcal{O}(N^3)$.

Theorem 2. [3] Let $M \in \mathbb{T}^{2N \times 2N}$ be a coherent DBM. Then M is octagonal-consistent if and only if M is consistent and $\lfloor \frac{M_{ii}}{2} \rfloor + \lfloor \frac{M_{jj}}{2} \rfloor \geq 0$, for all $1 \leq i, j \leq 2N$, $i \neq j$. Moreover, the tight closure of M is the DBM $M^t \in \mathbb{T}^{2N \times 2N}$ defined as $M_{ij}^t = \min \left\{ M_{ij}^*, \left\lfloor \frac{M_{ii}^*}{2} \right\rfloor + \left\lfloor \frac{M_{jj}^*}{2} \right\rfloor \right\}$, for all $1 \leq i, j \leq 2N$, where $M^* \in \mathbb{T}^{2N \times 2N}$ is the closure of M .

The tight closure of an octagonal constraint is needed for existential quantifier elimination, and ultimately, for proving that the class of octagonal relations is closed under composition [6].

Octagonal Relations are Ultimately Periodic Given a consistent octagonal relation $R(\mathbf{x}, \mathbf{x}')$ let $\sigma(R) = M_{\bar{R}}$. Dually, for any coherent DBM $M \in \mathbb{T}^{4N \times 4N}$, let $\rho(M) = \Omega_M$. Clearly, $\rho(\sigma(R)) \Leftrightarrow R$, as required by Def. 2.

In order to prove that the class \mathcal{R}_{oct} of octagonal relations is ultimately periodic, we need to prove that the sequence $\{\sigma(R^m)\}_{m=0}^{\infty}$ is ultimately periodic, for an arbitrary relation $R \in \mathcal{R}_{oct}$. It is sufficient to consider only the case where R is ω -consistent, hence $\sigma(R^m) = M_{\bar{R}^m}$, for all $m \geq 0$. We rely in the following on the main result of [6], which establishes a relation between $M_{\bar{R}^m}$ (the octagonal DBM corresponding to the m -th iteration of R) and $M_{\bar{R}^m}$ (the DBM corresponding to the m -th iteration of $\bar{R} \in \mathcal{R}_{db}$), for $m > 0$:

$$(M_{\bar{R}^m})_{ij} = \min \left\{ (M_{\bar{R}^m})_{ij}, \left\lfloor \frac{(M_{\bar{R}^m})_{i\bar{i}}}{2} \right\rfloor + \left\lfloor \frac{(M_{\bar{R}^m})_{\bar{j}j}}{2} \right\rfloor \right\}, \text{ for all } 1 \leq i, j \leq 4N \quad (*)$$

This relation is in fact a generalization of the tight closure expression from theorem 2, from $m = 1$ to any $m > 0$.

In Section 4.1 it was shown that difference bounds relations are ultimately periodic. In particular, this means that the sequence $\{M_{\bar{R}^m}\}_{m=0}^{\infty}$, corresponding to the iteration of the difference bounds relation \bar{R} , is ultimately periodic. To prove that the sequence $\{M_{\bar{R}^m}\}_{m=0}^{\infty}$ is also ultimately periodic, it is sufficient to show that: the minimum and the sum of two ultimately periodic sequences are ultimately periodic, and also that the integer half of an ultimately periodic sequence is also ultimately periodic.

Lemma 6. Let $\{s_m\}_{m=0}^{\infty}$ and $\{t_m\}_{m=0}^{\infty}$ be two ultimately periodic sequences. Then the sequences $\{\min(s_m, t_m)\}_{m=0}^{\infty}$, $\{s_m + t_m\}_{m=0}^{\infty}$ and $\{\lfloor \frac{s_m}{2} \rfloor\}_{m=0}^{\infty}$ are ultimately periodic as well.

Together with the above relation (*), lemma 6 proves that \mathcal{R}_{oct} is ultimately periodic.

Checking ω -Consistency and Inconsistency of Octagonal Relations This section describes an efficient method of deciding the queries \mathcal{Q}_1 and \mathcal{Q}_2 (lines 7 and 9 in Fig. 1) for the class of octagonal relations. In order to deal with these queries symbolically, we need to consider first the class $\mathcal{R}_{oct}(k)$ of octagonal relations with parameter k . In the rest of this section, let $k \notin \mathbf{x}$ be a variable ranging over \mathbb{N}_+ .

Definition 8. Then a formula $\phi(\mathbf{x}, z)$ is a parametric octagonal constraint if it is equivalent to a finite conjunction of terms of the form $\pm x_i \pm x_j \leq a_{ij} \cdot k + b_{ij}$, $2x_i \leq c_i \cdot k + d_i$, or $-2x_i \leq c'_i \cdot k + d'_i$, where $a_{ij}, b_{ij}, c_i, d_i, c'_i, d'_i \in \mathbb{Z}$ and $1 \leq i, j \leq N$, $i \neq j$.

A parametric octagon $\phi(\mathbf{x}, k)$ is represented by a matrix $M_{\overline{\phi}}[k] \in \mathbb{T}[k]^{2N \times 2N}$ of linear terms over k , and viceversa, a matrix $M[k] \in \mathbb{T}[k]^{2N \times 2N}$ corresponds to a parametric octagon $\Omega_M(k)$. We define $\pi(M[k]) = \Omega_M(k)$. As in the case of difference bounds constraints, one notices that $\pi(k \cdot \Lambda + \sigma(R^b)) \in \mathcal{R}_{oct}(k)$, for $R \in \mathcal{R}_{oct}$, $b \geq 0$ and $\Lambda \in \mathbb{T}^{4N \times 4N}$.

The composition of parametric octagonal relations (from e.g. \mathcal{Q}_1) requires the computation of the tight closure in the presence of parameters. According to theorem 2, the parametric tight closure can be expressed as a matrix of elements of the form $\min\{t_i(k)\}_{i=1}^m$, where $t_i(k)$ are either: (i) linear terms, i.e. $t_i(k) = a_i \cdot k + b_i$, or (ii) sums of halved linear terms, i.e. $t_i(k) = \lfloor \frac{a_i \cdot k + b_i}{2} \rfloor + \lfloor \frac{c_i \cdot k + d_i}{2} \rfloor$.

The main idea is to split a halved linear term of the form $\lfloor \frac{a \cdot k + b}{2} \rfloor$, $k > 0$ into two linear terms $a \cdot k + \lfloor \frac{b}{2} \rfloor$ and $a \cdot k + \lfloor \frac{b-a}{2} \rfloor$, corresponding to the cases of $k > 0$ being even or odd, respectively. This is justified by the following equivalence:

$$\{\lfloor \frac{a \cdot k + b}{2} \rfloor \mid k > 0\} = \{a \cdot k + \lfloor \frac{b}{2} \rfloor \mid k > 0\} \cup \{a \cdot k + \lfloor \frac{b-a}{2} \rfloor \mid k > 0\}$$

Hence, an expression of the form $\min\{t_i(k)\}_{i=1}^m$ yields two expressions $\min\{t_i^e(k)\}_{i=1}^m$, for even k , and $\min\{t_i^o(k)\}_{i=1}^m$, for odd k , where t_i^e and t_i^o , $1 \leq i \leq m$, are effectively computable linear terms. With these considerations, \mathcal{Q}_1 (for octagonal relations) is equivalent to a conjunction of equalities of the form $\forall k > 0. \min\{t_i^\bullet(k)\}_{i=1}^m = t_0^\bullet(k)$, $\bullet \in \{e, o\}$. Now we can apply lemma 4 to the right-hand sides of the equivalences above, to give efficient equivalent conditions for deciding \mathcal{Q}_1 .

The query \mathcal{Q}_2 is, according to theorem 2, equivalent to finding either (i) a strictly negative cycle in a parametric octagonal DBM $M[k]$, or (ii) a pair of indices $1 \leq i, j \leq 4N$, $i \neq j$ such that $\lfloor \frac{M[k]_{ii}}{2} \rfloor + \lfloor \frac{M[k]_{jj}}{2} \rfloor < 0$. Considering that the set of terms corresponding to the two cases above is $T = \{a_i \cdot k + b_i\}_{i=1}^m \cup \{\lfloor \frac{c_i \cdot k + d_i}{2} \rfloor + \lfloor \frac{e_i \cdot k + f_i}{2} \rfloor\}_{i=1}^p$, we split each term $t \in T$ into two matching linear terms, and obtain, equivalently:

$$T_{e,o} = \{\alpha_i^e \cdot k + \beta_i^e\}_{i=1}^{m+p} \cup \{\alpha_i^o \cdot k + \beta_i^o\}_{i=1}^{m+p}$$

Now we can apply lemma 5, and compute the minimal value for which a term $t \in T_{e,o}$ becomes negative, i.e. $n_0 = \min_{i=1}^{m+p} \min(2\gamma_i^e, 2\gamma_i^o - 1)$, where $\gamma_i^\bullet = \max(1, \lfloor -\frac{\beta_i^\bullet}{\alpha_i^\bullet} \rfloor + 1)$, if $\alpha_i^\bullet < 0$, 1 if $\alpha_i^\bullet \geq 0 \wedge \alpha_i^\bullet + \beta_i^\bullet < 0$, and ∞ , otherwise, for $\bullet \in \{e, o\}$.

4.3 Finite Monoid Affine Transformations

The class of affine transformations is one of the most general classes of deterministic transition relations involving integer variables. If $\mathbf{x} = \langle x_1, \dots, x_N \rangle$ is a vector of variables ranging over \mathbb{Z} , an *affine transformation* is a relation of the form:

$$T \equiv \mathbf{x}' = A \otimes \mathbf{x} + \mathbf{b} \wedge \phi(\mathbf{x}) \quad (4)$$

where $A \in \mathbb{Z}^{N \times N}$, $\mathbf{b} \in \mathbb{Z}^N$, ϕ is a Presburger formula, and \otimes stands for the standard matrix multiplication in \mathbb{Z} .

The affine transformation is said to have the *finite monoid property* [5, 10] if the monoid $\langle \mathcal{M}_A, \otimes \rangle$, where $\mathcal{M}_A = \{A^{\otimes i} \mid i \geq 0\}$ is finite. In this case, we also say that A is finite monoid. Here $A^{\otimes 0} = I_N$ and $A^{\otimes i} = A \otimes A^{\otimes i-1}$, for $i > 0$. Intuitively, the finite monoid property is equivalent to the fact that A has finitely many powers (for

the standard integer multiplication) that repeat periodically. It is easy to see that A is finite monoid if and only if there exists $p \geq 0$ and $l > 0$ such that $A^{\otimes p} = A^{\otimes p+l}$, i.e. $\mathcal{M}_A = \{A^{\otimes 0}, \dots, A^{\otimes p}, \dots, A^{\otimes p+l-1}\}$.

If A is finite monoid, it can be shown that T^* can be defined in Presburger arithmetic [5, 10]. We achieve the same result below, by showing that finite monoid affine transformations are ultimately periodic relations. As a byproduct, the transitive closure of such relations can also be computed by the algorithm in Fig. 1.

An affine transformation T (4) can be equivalently written in the homogeneous form:

$$T \equiv \mathbf{x}'_h = A_h \otimes \mathbf{x}_h \wedge \phi_h(\mathbf{x}_h) \text{ where } A_h \equiv \left(\begin{array}{c|c} A & \mathbf{b} \\ \hline 0 \dots 0 & 1 \end{array} \right)$$

where $x_h = \langle x_1, \dots, x_N, x_{N+1} \rangle$ with $x_{N+1} \notin \mathbf{x}$ being a fresh variable and $\phi_h(\mathbf{x}_h) \equiv \phi(\mathbf{x}) \wedge x_{N+1} = 1$. In general, the k -th iteration of an affine transformation can be expressed as:

$$T^k \equiv \mathbf{x}'_h = A_h^{\otimes k} \otimes \mathbf{x}_h \wedge \forall 0 \leq \ell < k. \phi_h(A_h^{\otimes \ell} \otimes \mathbf{x}_h) \quad (5)$$

Notice that, if $\mathbf{x}_h^{(0)}$ denotes the initial values of \mathbf{x}_h , the values of \mathbf{x}_h at the ℓ -th iteration are $\mathbf{x}_h^{(\ell)} = A_h^{\otimes \ell} \otimes \mathbf{x}_h^{(0)}$. Moreover, we need to ensure that all guards up to (and including) the $(k-1)$ -th step are satisfied, i.e. $\phi_h(A_h^{\otimes \ell} \otimes \mathbf{x}_h)$, for all $0 \leq \ell < k$.

For the rest of the section we fix A and \mathbf{b} , as in (4). The encoding of a consistent affine transformation T is defined as $\sigma(T) = A_h \in \mathbb{T}^{(N+1) \times (N+1)}$. Dually, for some $M \in \mathbb{T}[k]^{(N+1) \times (N+1)}$, we define:

$$\pi(M) : \exists x_{N+1}, x'_{N+1}. \mathbf{x}'_h = M \otimes \mathbf{x}_h \wedge \forall 0 \leq \ell < k. \phi_h(M[\ell/k] \otimes \mathbf{x}_h)$$

where $M[\ell/k]$ denotes the matrix M in which each occurrence of k is replaced by ℓ . In contrast with the previous cases (Section 4.1 and Section 4.2), only M is not sufficient here to recover the relation $\pi(M) - \phi$ needs to be remembered as well⁶.

With these definitions, we have $\sigma(T^k) = A_h^{\otimes k}$, for all $k > 0$ – as an immediate consequence of (5). The next lemma proves that the class of finite monoid affine relations is ultimately periodic.

Lemma 7. *Given a finite monoid matrix $A \in \mathbb{Z}^{N \times N}$ and a vector $\mathbf{b} \in \mathbb{Z}^N$, the sequence $\{A_h^{\otimes k}\}_{k=0}^{\infty}$ is ultimately periodic.*

The queries \mathcal{Q}_1 and \mathcal{Q}_2 (lines 7 and 9 in Fig. 1) in the case of finite monoid affine transformations, are expressible in Presburger arithmetic. These problems could be simplified in the case of affine transformations *without guards*, i.e. $T \equiv \mathbf{x}' = A\mathbf{x} + \mathbf{b}$. The transformation is, in this case, ω -consistent. Consequently, \mathcal{Q}_1 reduces to an equivalence between two homogeneous systems $\mathbf{x}'_h = A_{1h} \otimes \mathbf{x}_h$ and $\mathbf{x}'_h = A_{2h} \otimes \mathbf{x}_h$. This is true if and only if $A_{1h} = A_{2h}$. The query \mathcal{Q}_2 becomes trivially false in this case.

⁶ This incurs a slight modification of the algorithm presented in Fig. 1.

5 Experimental Results

We have implemented the transitive closure algorithm from Fig. 1 within the FLATA toolset [11], a framework we develop for the analysis of counter systems. We compared the performance of this algorithm with our older transitive closure computation methods for difference bounds [8] and octagonal relations [6]. We currently lack experimental data for finite monoid relations (namely, a comparison with existing tools such as FAST [4], LASH [14] or TRex [2] on this class), as our implementation of finite monoid affine transformation class is still underway.

Table 1 shows the results of the comparison between the older algorithms described in [8, 6] (denoted as **old**) and the algorithm in Fig. 1 for difference bounds relations $d_{1,\dots,6}$ and octagonal relations $o_{1,\dots,6}$. The tests have been performed on both **compact** (minimum number of constraints) and **canonical** (i.e. closed, for difference bounds and tightly closed, for octagons) relations. The **speedup** column gives the ratio between the **old** and **new** execution times. The experiments were performed on a 2.53GHz machine with 2.9GB of memory.

Relation	new	compact		canonical	
		old	speedup	old	speedup
d_0 $(x - x' = -1) \wedge (x = y')$	0.18	0.7	3.89	38.77	215.39
d_1 $(x - x' = -1) \wedge (x' = y')$	0.18	18.18	101.0	38.77	215.39
d_2 $(x - x' = -1) \wedge (x = y') \wedge (x - z' \leq 5) \wedge (z = z')$	1.2	26.5	22.1	33431.2	27859.3
d_3 $(x - x' = -1) \wedge (x = y') \wedge (x - z \leq 5) \wedge (z = z')$	0.6	32.7	54.5	33505.5	55841.7
d_4 $(x - x' = -1) \wedge (x = y) \wedge (x - z \leq 5) \wedge (z = z')$	0.5	702.3	1404.6	48913.8	97827.6
d_5 $(a = c) \wedge (b = a') \wedge (b = b') \wedge (c = c')$	1.8	5556.6	3087.0	$> 10^6$	∞
d_6 $(a - b' \leq -1) \wedge (a - e' \leq -2) \wedge (b - a' \leq -2)$ $\wedge (b - c' \leq -1) \wedge (c - b' \leq -2) \wedge (c - d' \leq -1)$ $\wedge (d - c' \leq -2) \wedge (d - e' \leq -1) \wedge (e - a' \leq -1)$ $\wedge (e - d' \leq -2) \wedge (a' - b \leq 4) \wedge (a' - c \leq 3)$ $\wedge (b' - c \leq 4 \wedge b' - d \leq 3) \wedge (c' - d \leq 4) \wedge (c' - e \leq 3)$ $\wedge (d' - a \leq 3 \wedge d' - e \leq 4) \wedge (e' - a \leq 4) \wedge (e' - b \leq 3)$	5.6	$> 10^6$	∞	$> 10^6$	∞
o_1 $(x + x' = 1)$	0.21	0.91	4.33	0.91	4.33
o_2 $(x + y' \leq -1) \wedge (-y - x' \leq -2)$	0.29	0.85	2.93	0.84	2.9
o_3 $(x \leq x') \wedge (x + y' \leq -1) \wedge (-y - x' \leq -2)$	0.32	0.93	2.91	0.94	2.94
o_4 $(x + y \leq 5) \wedge (-x + x' \leq -2) \wedge (-y + y' \leq -3)$	0.21	3.67	17.48	13.52	64.38
o_5 $(x + y \leq 1) \wedge (-x \leq 0) \wedge (-y \leq 0)$	1.2	20050.9	16709.1	$> 10^6$	∞
o_6 $(x \geq 0) \wedge (y \geq 0) \wedge (x' \geq 0) \wedge (y' \geq 0)$ $\wedge (x + y \leq 1) \wedge (x' + y' \leq 1) \wedge (x - 1 \leq x')$ $\wedge (x' \leq x + 1) \wedge (y - 1 \leq y') \wedge (y' \leq y + 1)$	2.5	$> 10^6$	∞	$> 10^6$	∞

Table 1: Comparison with older algorithms on difference bounds and octagons. Times are in milliseconds.

As shown in Table 1, the maximum observed speedup is almost 10^5 for difference bounds (d_4 in canonical form) and of the order of four for octagons. For the relations d_5 (canonical form), d_6 and o_6 the computation using older methods took longer than 10^6 msec. It is also worth noticing that the highest execution time with the new method was of 2.5 msec.

Table 2 compares FLATA with FAST [4] on counter systems with one self loop labeled with a randomly generated deterministic difference bound relation. We generated 50 such relations for each size $N = 10, 15, 20, 25, 50, 100$. Notice that FAST usually runs out of memory for more than 25 variables, whereas FLATA can handle 100 variables in reasonable time (less than 8 seconds on average).

vars	FLATA			FAST				
	done	av.	E_T	done	av.	E_T	E_M	E_B
10	50	1.5	0	49	0.6	0	0	1
15	50	1.6	0	31	10.5	17	0	2
20	50	1.6	0	4	3.4	9	8	29
25	50	1.6	0	2	4.2	2	10	36
50	50	1.6	0	0	-	0	0	50
100	49	7.7	1	0	-	0	0	50

(a) – matrix density 3%

vars	FLATA			FAST				
	done	av.	E_T	done	av.	E_T	E_M	E_B
10	50	1.5	0	22	6.9	23	1	4
15	50	1.5	0	1	20.6	4	3	42
20	50	1.6	0	0	-	1	0	49
25	43	1.7	7	0	-	0	0	50
50	50	2.3	0	0	-	0	0	50
100	42	5.5	8	0	-	0	0	50

(b) – matrix density 10%

Table 2: Comparison with FAST (MONA plugin) on deterministic difference bounds. Times are in seconds. E_T – timeout 30 s, E_B – BDD too large, E_M – out of memory

6 Conclusion

We presented a new, scalable algorithm for computing the transitive closure of ultimately periodic relations. We show that this algorithm is applicable to difference bounds, octagonal and finite monoid affine relations, as all three classes are shown to be ultimately periodic. Experimental results show great improvement in the time needed to compute transitive closures of difference bounds and octagonal relations.

References

1. A. Annichini, E. Asarin, and A. Bouajjani. Symbolic techniques for parametric reasoning about counter and clock systems. In *CAV '00*, pages 419–434. Springer, 2000.
2. A. Annichini, A. Bouajjani, and M. Sighireanu. Trex: A tool for reachability analysis of complex systems. In *CAV '01*, pages 368–372. Springer, 2001.
3. R. Bagnara, P. M. Hill, and E. Zaffanella. An improved tight closure algorithm for integer octagonal constraints. In *VMCAI '08*, 2008.
4. S. Bardin, J. Leroux, and G. Point. Fast extended release. In *CAV'06*, volume 4144, pages 63–66. Springer Verlag, 2006.
5. B. Boigelot. *Symbolic Methods for Exploring Infinite State Spaces*, volume PhD Thesis, Vol. 189. Collection des Publications de l'Université de Liège, 1999.
6. M. Bozga, C. Gîrlea, and R. Iosif. Iterating octagons. In *TACAS '09*, pages 337–351. Springer, 2009.
7. M. Bozga, R. Iosif, and F. Konečný. Fast Acceleration of Ultimately Periodic Relations. Technical Report TR-2010-3, Verimag, Grenoble, France, 2010.
8. M. Bozga, R. Iosif, and Y. Lakhnech. Flat parametric counter automata. *Fundamenta Informaticae*, 91:275–303, 2009.
9. H. Comon and Y. Jurski. Multiple Counters Automata, Safety Analysis and Presburger Arithmetic. In *CAV '98*, volume 1427 of *LNCS*, pages 268 – 279. Springer, 1998.
10. A. Finkel and J. Leroux. How to compose presburger-accelerations: Applications to broadcast protocols. In *FST TCS '02*, pages 145–156. Springer, 2002.
11. <http://www-verinew.imag.fr/FLATA.html>.
12. A. Miné. The octagon abstract domain. *Higher-Order and Symbolic Computation*, 19(1):31–100, 2006.
13. B. De Schutter. On the ultimate behavior of the sequence of consecutive powers of a matrix in the max-plus algebra. *Linear Algebra and its Applications*, 307:103–117, 2000.
14. P. Wolper and B. Boigelot. Verifying systems with infinite but regular state spaces. In *CAV '98*, pages 88–97. Springer, 1998.