

A Modular Static Analysis Approach to Affine Loop Invariants Detection

Corinne Ancourt^{1,2}, Fabien Coelho³ and François Irigoin⁴

*CRI, Maths & Systems, MINES ParisTech,
35, rue Saint Honoré, 77300 Fontainebleau, France*

Abstract

Modular static analyzers use procedure abstractions, a.k.a. summarizations, to ensure that their execution time increases linearly with the size of analyzed programs. A similar abstraction mechanism is also used within a procedure to perform a bottom-up analysis. For instance, a sequence of instructions is abstracted by combining the abstractions of its components, or a loop is abstracted using the abstraction of its loop body: fixed point iterations for a loop can be replaced by a direct computation of the transitive closure of the loop body abstraction.

More specifically, our abstraction mechanism uses affine constraints, i.e. polyhedra, to specify pre- and post-conditions as well as state transformers. We present an algorithm to compute the transitive closure of such a state transformer, and we illustrate its performance on various examples. Our algorithm is simple, based on discrete differentiation and integration: it is very different from the usual abstract interpretation fixed point computation based on widening. Experiments are carried out using previously published examples. We obtain the same results directly, without using any heuristic.

Keywords: Abstract interpretation, fixed point computation, loop invariant.

1 Introduction

Program analyses such as interprocedural program parallelization [21,20], array access bound checking [26], array initialization checking, aliasing checking [25] require some mechanism to approximate loop behaviors. In order to obtain a modular analyzer and to limit analysis times, we depart from the usual approach [8] and compute state transformers instead of state predicates, *i.e.* pre- and post-conditions. Transformers are used to summarize functions: each function is analyzed once and its transformer is reused at each call site. Preconditions are then propagated using the transformers. Since transformers require twice as many variables as preconditions,

¹ Work funded by the ACI SI as part of APRON – <http://www.cri.enscm.fr/apron>

² Email: Corinne.Ancourt@cri.mines-paristech.fr

³ Email: Fabien.Coelho@cri.mines-paristech.fr

⁴ Email: Francois.Irigoin@cri.mines-paristech.fr

we use polyhedra as finite abstractions of possibly infinite sets of states to maintain a sufficient accuracy.

In Section 2, we present a simple yet effective algorithm to compute transitive closures of transformers, which are then used to derive affine loop invariants. Then we show how to improve its effectiveness by using equivalent but different formulae for postconditions. They are equivalent when the analysis is exact, but they differ when approximations, such as affine approximations, are made. Several kinds of extensions are considered in Section 3. They are related to the transitive closure algorithm. Related work is introduced in Section 4 and we show on previously published examples that our algorithm provides the expected loop invariants without using any widening heuristic.

2 Simple Transitive Closure of Affine Transformers

Pugh and al. studied the transitive closure of transfer functions defined by Presburger formulae [22]. Here, transfer functions are approximated by affine relations. The graph of the relation between the initial state and the final state is defined by a polyhedron, i.e. a set of affine equalities and inequalities. Below, once transformers and preconditions are defined, we present our algorithm to compute transformer closures for while loops together with its proof. We illustrate its working on a motivating example, a safety controller for a toy robot car [17].

2.1 Affine Transformers and Preconditions

Each program command, elementary or compound statement or procedure call is approximated by an affine transformer. The underlying mechanism is similar to [8] but extended from the states to state transitions. The idea of transformers is quite general and is also used, for instance by Boigelot & al. [4].

The set of possible program states, before a command is executed, is defined by a precondition. The set of program states after the command execution is defined by a postcondition. The postcondition is the image of the precondition by the command transformer. A legal affine abstract postcondition contains the effective postcondition, i.e. it is an over-approximation.

For simplicity of exposure, the relationship between identifiers and memory locations is assumed to be a one-to-one mapping for scalar variables. In this paper, we deal only with integer scalar variables and states taking values in \mathbb{Z}^n , where the dimension n is the number of analyzed variables.

For semantic analysis purposes, control flow graphs are structured as while loops, e.g. using Bourdoncle's heuristic [5]. Other structured loops are decomposed into while loops. As a result, the only control structure with an iterative behavior studied here is the simple while loop.

2.2 The Affine Derivative Closure Algorithm

Our algorithm is outlined in Fig. 1 and works as follows:

```

transformer T*(x,x') affine_derivative_closure(transformer T(x,x'))
{
  // add the difference vector dx
  transformer Q(x,x',dx) = T(x,x') ^ (dx = x'-x);
  // eliminate the initial and final states x and x'
  transformer T'(dx) = project((project(Q(x,x',dx), x), x'));
  // compute dx for any iteration number k
  T'(dx) = multiply_constant_terms(T'(dx), k) ^ (k >= 0);
  // eliminate the iteration number k and substitute back dx by x'-x
  return project(project(T'(dx),k) ^ (dx = x'-x), dx);
}

```

Fig. 1. Affine Derivative Closure Algorithm

Let us assume that T is a valid affine transformer for a while loop body and its continuation condition. T includes the loop entry condition, or at least an affine approximation of this condition. Let k be the iteration number, x^{k-1} be the integer memory state when the loop body is started the k -th time, and x^k be the final state when the loop condition is evaluated to true again. The predicate $T(x^{k-1}, x^k)$ holds for all possible $k > 1$.

Let δx be $x^k - x^{k-1}$ and $T'(\delta x)$ be the projection of $T \wedge \delta x = x^k - x^{k-1}$ along x^k and x^{k-1} . Note that T' does not depend neither on k which is not a component of the memory state x nor on the names x^k and x^{k-1} , which have been eliminated by the projection.

Let x^0 be the state on loop entry. The state x^k that may be reached after k iterations of the loop, if such an iteration is executed, is:

$$x^k = x^0 + \sum_{i=1}^k \delta x^i \quad \text{with } \delta x^i = x^i - x^{i-1} \quad (1)$$

For all positive integers i , $T'(\delta x^i)$ holds. Since T' is a polyhedron, it can be defined by affine equalities and inequalities:

$$T'(\delta x) = \{ \delta x \mid A\delta x = b \wedge A'\delta x \leq b' \} \quad (2)$$

where A and A' are integer matrices and b and b' the corresponding constant terms.

Multiplying Eq.(1) by the matrices A and A' , we have:

$$Ax^k = Ax^0 + \sum_{i=1}^k A\delta x^i \quad \text{and} \quad A'x^k = A'x^0 + \sum_{i=1}^k A'\delta x^i \quad (3)$$

Since $A\delta x^i$ and $A'\delta x^i$ are equal to b or bounded by b' , we have:

$$Ax^k = Ax^0 + kb \quad \text{and} \quad A'x^k \leq A'x^0 + kb' \quad (4)$$

The loop transformer $T^*(x^0, x)$ may be approximated by:

$$T^*(x^0, x) \subseteq \{ (x^0, x) \mid \exists k \in [0, \infty[\quad Ax = Ax^0 + kb \wedge A'x \leq A'x^0 + kb' \} \quad (5)$$

The set $P^*(x^0)$ of states reachable from x^0 by executing the loop may thus be approximated by:

$$P^*(x^0) \subseteq \{ x \mid T^*(x^0, x) \} \quad (6)$$

Then the whole loop transformer $T^*(x^0, x)$ is over-approximated by projecting k from the constraints in Eq.(5). And an affine over-approximation of the loop

```

1:  int s = 0, t = 0, d = 0;
2:  while(s <= 2 && t <= 3)
3:    if(alea())
4:      t++, s = 0; // increment time, reset speed estimation
5:    else
6:      d++, s++; // meter increment, speed estimation increment

```

Fig. 2. Car safety example

postcondition is obtained by projecting x^0 too and by adding a safe approximation of the loop last iteration and exit condition.

Loop invariants are obtained for each state dimension i such that b_i is zero since then Eq.(4) shows that $(Ax^k)_i = (Ax^0)_i$. If there exists a dimension j such that b_j is not zero, then the iteration number k can be derived from a combination of variables and substituted everywhere else to obtain more invariants.

If not, k is still bounded by $k \geq 0$ and some inequalities can be saved according to the Fourier-Motzkin elimination rule. If some term $a'\delta x$, where a' is a row of A' , is upper-bounded by a negative constant, or lower-bounded by a positive constant, monotony constraints are obtained. Strict monotonicity leads to loop termination proofs when the derivatives of the affine components of the while condition that imply non-termination are incompatible with T' .

Note that $T^+ = T^* \circ T$ can be computed by setting $k \geq 1$ in Eq.(5). Transformer T^+ may contain strict monotonicity conditions, which are useful for dependence testing in automatic loop parallelization [30] and array bound checking [29].

2.3 An Example: Robot Car Safety

Let us take the toy example described in [17] and recently reused by [23]. A robot car must follow autonomously a track painted on the floor. In case it loses the track, it should not crash against a wall; however it is not stopped right away since the track might be found again. The car should not accelerate too much when it is looking for the lost track. The safety controller must ensure that a limited amount of time is allowed to search the painted track at bounded speed. Since time and speed are bounded in the track search mode, the car is safe if the track is far enough from the walls.

Let t be the time in seconds, d the distance from the starting point in meters and s the current estimation of the speed in meters per second. A model of the controller ensuring the physical safety of the car is encoded in C as shown in Fig. 2. Function *alea* is used to model a random event: either the clock counter is going to tick for the next second and the time is incremented while the speed estimation is reset to 0, or the distance and the speed estimation are increased because another meter has been reached. The safety is enforced by the loop guard. If nothing else happens within three seconds or if the speed is greater than two meters per second, the car is stopped. If walls are 10 meters away from the starting position, the car cannot reach a wall.

We explain the steps performed here by our Affine Derivative Closure Algorithm using a primed notation that distinguishes the values of each variable between the old and primed new state. If x is the memory state of (d, s, t) , the transformer for

the first branch of the test (line 4) is:

$$T_4(x, x') = \{s' = 0, t' = t + 1, d' = d, s \leq 2, t \leq 3\}$$

For the second test branch (line 6) the transformer is:

$$T_6(x, x') = \{d' = d + 1, s' = s + 1, s \leq 2, t' = t, t \leq 3\}$$

Their convex hull used to approximate their combined effect is:

$$T_3(x, x') = \{d' + t' = d + t + 1, s + 3t + 1 \leq s' + 3t' \leq 3t + 3, t \leq t' \leq t + 1, t \leq 3\}$$

Projecting the old and new state, this transformer is rewritten as:

$$T'(\delta x) = \{\delta d + \delta t = 1, 1 \leq \delta s + 3\delta t, 0 \leq \delta t \leq 1\}$$

which leads to $\delta d + \delta t \leq \delta s + 3\delta t$, or $\delta d \leq \delta s + 2\delta t$. This is the speed equation we looked for to prove the car safety. If the speed and the time are bounded, the distance travelled is bounded. Here, the numerical speed bound of 2 produces a linear speed equation.

2.4 Discussion

The algorithm is very simple yet powerful enough to derive non trivial conditions. Its weaknesses come from 1) computing T' as a relation on δx instead of a more accurate relation on $(\delta x, x)$ to ease the summation and stay in the affine setting, and 2) in computing T in the first place as an affine transformer using the convex hull to model tests. The complexity of the algorithm is dominated by the complexity of the projection steps: its worst case is exponential with the number of variables projected, but in practice it is polynomial when the constraints are sparse.

3 From Transformers to Loop Invariants

Several simple extensions are useful to cope with non-affine behaviors such as iteration independent assignments or periodic and polynomial behaviors. They occur when an iteration independent assignment is equivalent to a differential assignment, when two (or more) buffers are used in a flip-flop mode or when triangular matrices are accessed. Note that polynomial behaviors [15] are frequent when accessing symmetric matrices, but that monotonicity or strict monotonicity information is often sufficient to make a decision about data dependence or array bound overflow issues.

3.1 Using T^+ instead of T^*

If the loop w is certainly entered when reached with precondition P^w , that is when the affine approximation of the negation of its condition combined with P^w generates a contradiction, it is better to compute T^+ instead of T^* . The constraints on the image of T can be added to the image of T^+ , but not of T^* .

In a loop such as “while(alea()) m = 10;” no information on m is gathered in the postcondition because its value may be unchanged, when the loop is not entered, as well as set to 10.

Note that $T^+ = T \circ T^* = T^* \circ T$ when T is exact, but that the first formula is more precise with an approximate T . In the second case, the information added by

T may be lost by T^* . Also it is better to use:

$$P^* = P^0 \sqcup T^+(P^0) \quad (7)$$

rather than the equivalent formula $P^* = T^*(P^0)$ when the range of T and P^0 have common constraints. The convex hull operator \sqcup is used instead of the union operator, which is not internal for polyhedra. Since it is not accurate, it should always be applied as late as possible when equivalent formulae are available.

3.2 Periodic Behaviors

Periodic behaviors are observed when a variable or a set of variables is used to flip-flop the accesses to two or more buffers; for instance this is often used in signal processing applications to switch between receiving or sending and computing buffers, or in scientific programs to switch between new and old values [4,3], as depicted in Fig. 3.

```
double x[2][10];
int old = 0, new = 1, i, t;
for(t = 0; t<1000; t++) {
  for(i = 0; i<10;i++)
    x[new][i] = g(x[old][i]);
  old = new, new = 1 - old;
}
```

Fig. 3. Flip-flop example

The t loop is parallel if the value of new is proven to be always different from the value of old because $old + new = 1$, which is found thanks to Eq.(7).

An interesting aspect in flip-flop analysis is its robustness with respect to the flip-flop encoding scheme. Ideally, different encodings leading to the same execution traces should produce the same analysis result. Different encodings use different mathematical functions, as illustrated by Figure 4, but they have the same value over the useful subset of their domains. Thus the analysis result depends on the accuracy of the loop precondition used to characterize the effective function domains.

1	<code>new = old; old = 1-old;</code>
2	<code>new = 1 - new; old = 1 - old;</code>
3	<code>t = new; new = old; old = t;</code>
4	<code>if(new==1) { new = 0, old = 1; } else { new = 1, old = 0; }</code>
5	<code>if(new==1) { new = 0; old = 1; } else { if(new==0) { new = 1; old = 0; } else exit(1); }</code>
6	<code>new = (new+1) %2; old = (old+1) %2 ;</code>

Fig. 4. Six different encodings of **flip-flop** operations

The invariant $new + old = 1$ is found by our tool PIPS [24] for Cases 1, 3, 4 and 5 thanks to Eq.(7). But it fails for Case 6 because the modulo operator is not

analyzed as well as a multiplication or a division: the sources of failure are not limited to convex hulls and transitive closures.

Case 2 requires a loop unrolling of two to obtain the invariant. Larger periods can be obtained using integer rotation matrices or ad hoc constructs. More generally, k -periodic behaviors can be captured by computing T^* and T^+ as:

$$T_k^* = \bigsqcup_{i=0, k-1} T^i \circ (T^k)^* \quad T_k^+ = \bigsqcup_{i=1, k} T^i \circ (T^k)^* \quad (8)$$

This is equivalent to a loop unrolling of degree k and similar to a delayed widening. These definitions can be used to refine Eq.(7) and to obtain better loop preconditions. With $k = 2$, the precondition becomes:

$$P^* = P^0 \bigsqcup T(P^0) \bigsqcup T^{2+}(P^0) \bigsqcup T(T^{2+}(P^0)) \quad (9)$$

With Eq.(9), PIPS is able to deal with the second encoding of flip-flop because `new` and `old` are invariant by T^2 . The effective period does not have to be known as each T_k^* is a proper over-approximation of T^* and their intersection can be used: $T^* = \bigcap_{i \in [0, k[} T_i^*$. The same holds for T^+ .

Theorem 3.1 (Invariant Completeness Theorem) *Any invariant found by our Affine Derivative Closure Algorithm for transformer T^* is also found by the same algorithm for transformer $(T^n)^*$.*

The proof is in the Appendix. In other words, the accuracy can only be improved when n is increased.

3.3 Higher-Order Differences

The scheme could be first generalized to second order differences by setting: $T(x, x') \wedge \delta x = x' - x \wedge T(x', x'') \wedge \delta x' = x'' - x' \wedge \Delta x = \delta x' - \delta x$

Let us consider the non linear example on the left hand side of Fig. 5. One possible application of such differences is to prove that variable `i` is bounded and reaches its maximum on the loop boundary or when its discrete difference is zero. Indeed since the second difference of `i` is the difference of `j`, -1 , and is negative, sooner or later, `i` is going to decrease.

<pre>int i = 0, j = 2, k = 1; while(k<=10) j--, i += j, k++;</pre>	<pre>int i = 0, j = 0, n; if(n<0) exit(1); while(i<=n) i++, j+=i;</pre>
---	---

Fig. 5. Non linear (left) and parabolic (right) examples

Exact closed form polynomials are computed in [15] and [28], but the polynomial closed form would be uselessly complicated to use for this purpose, although admittedly mandatory for code generation after automatic parallelization.

3.4 Monotonicity and Iterative Analysis

Postcondition $\{j = -8, k = 11\}$ is directly derived from the code in the left of Fig. 5. It does not bound `i`. However, if the transformers are recomputed with this

precondition, the loop postcondition is refined iteratively as shown on the left hand side of Fig 6.

A similar code, without loop numerical bound, on the right hand side in Fig. 5, can also be analyzed iteratively, but without ever reaching a fixed point, as shown on the right hand side of Fig 6.

<pre> 1 {j = -8, k = 11, 0 ≤ i + 71, i + 14 ≤ 0} 2 {j = -8, k = 11, 0 ≤ i + 71, i + 25 ≤ 0} 3 {j = -8, k = 11, 0 ≤ i + 71, i + 32 ≤ 0} 4 {j = -8, k = 11, 0 ≤ i + 71, i + 35 ≤ 0} 5 {j = -8, k = 11, 0 ≤ i + 71, i + 35 ≤ 0} </pre>		<pre> 1 {i = n + 1, 1 ≤ i} 2 {..., 2i ≤ j + 1, 3i ≤ j + 3, 4i ≤ j + 6} 3 {..., 5i ≤ j + 10, 6i ≤ j + 15} 4 {..., 7i ≤ j + 21, 8i ≤ j + 28} 5 {..., 9i ≤ j + 36, 10i ≤ j + 45} </pre>
--	--	---

Fig. 6. Results for the non linear and parabolic examples of Fig. 5

The iterative relationship between transformers and preconditions is formalized by the next two equations where B stands for the loop body statement and the continuation condition, and \mathcal{T} for the function that converts a statement into a convex transformer:

$$T_{n+1}^* = \mathcal{T}(B, P_n^*) \wedge P_n^* \quad P_n^* = P^0 \sqcup T_n(T_n^*(P^0)) \quad (10)$$

Note that the previous precondition P_n impacts the transformer T_{n+1} in two different ways. The affine abstraction \mathcal{T} is sharpened and the resulting transformer also is restricted by the previous precondition.

The iterative refinement process does not always converge and it may even lead to a precision loss, due to magnitude overflows. These are not handled with a good heuristic in the present PIPS implementation, but it is not critical as the refinement process is not automatic: it must be specified by the user.

3.5 Postponing Convex Hulls

If a loop contains a test, the test is abstracted by a convex hull and the transitive closure is applied later. In other words, the convex hull loses information at the very beginning of the invariant computation.

Hence it is useful to convert: `while(c) if(t) a; else b;`

into the equivalent: `while (c) { while (c&& t) a; while (c&&!t) b; }`

This transformation, which is somehow similar to the *abstract acceleration* defined by Laure Gonnord [10] after the acceleration used in model-checking [4,6,7], eliminates the early convex hull and lets PIPS find the proper invariants for cases 1 in [11] and 2 in [14,13].

4 Related Work

PIPS [24] development has been driven for almost twenty years by the needs of automatic analysis and parallelization for large size real-life Fortran and C programs of up to hundreds of functions and 100 KLOCs. Our derivative algorithm is used for all loops and on most of the control-flow graphs, after restructuring.

Original example	After restructuring	PIPS
(1) <i>Gopan'07 [11], Gulwani'09 [13]</i> <pre>x=y=0; while(*) { if (x ≤ 50) y++; else y--; if (y < 0) break; x++; }</pre>	<pre>x=y=0; while (y ≥ 0) { while (y ≥ 0 && x ≤ 50) y++; x++; while (y ≥ 0 && x > 50) y--; x++; } x--;</pre>	$x = 102$
(2) <i>Gulwani 2007 [14,13]</i> <pre>x=1; y=50; while(x < 100) { if (x < 50) x++; else { x++; y++; } }</pre>	<pre>x=1; y=50; while (x < 100) { while (x < 50) x++; while (x < 100 && x ≥ 50) { x++; y++; } }</pre>	$y = 100$
(3) <i>Gulavani 2006 [12], Gulwani 2009 [13]</i> <pre>x=1; lock=y=0; while (x ≠ y) { lock=1; x=y; if (*) lock =0; y++; }</pre>		lock = 1
(4) <i>Gaz Burner - Chaochen [31], Gonnord [10]</i> <pre>t=l=x=0; while(*) { x=0; while (x ≤ 9 && alea()) x++, t++, l++; x=0; while (x ≤ 49 alea()) x++, t++;}</pre>		$6l \leq t + 5x$
(5) <i>Halbwachs [16]</i> <pre>x=y=0; while(x ≤ 100) { if (alea()) x = x+2; else x++, y++; }</pre>		$2 \leq x + y$ $y \leq x$ $x + y \leq 202$
(6) <i>Halbwachs [16]</i> <pre>x=y=0; while (x ≤ 100) { if (alea()) x = x+4; else x=x+2, y++; }</pre>		$4 \leq x + 2y$ $2y \leq x$ $x + 2y \leq 204$
(7) <i>Robot car safety example</i> <i>Halbwachs [17], Merchat [23], Gonnord [10], Sect. 2.3-Fig. 2</i>		$d \leq s + 2t$
(8) <i>Subway example</i> <i>Halbwachs [19], Gonnord [10]</i>		$20 \leq b$ $b - s \leq 20$

Fig. 7. PIPS Experimental results

Its input is a deterministic C or Fortran program and not a non-deterministic finite automata as used in model-checking benchmarks or by Gonnord [10]. It is difficult to be sure to convert an automaton into a program without performing some intelligent structuring that may turn out to be the key to its successful analysis. Moreover, the PIPS semantic analyzer uses Bourdoncle's algorithm [5] to deal with unstructured control flow graph. This increases the number of convex hulls to decrease the number of widenings, and we now know that it often prevents acceleration opportunities. Most of our comparisons are based on published pieces of code, not on transcoded automata.

Fig. 7 presents examples found in the literature [8,10,11,12,14,16,17,23,13] about the widening operator and its improvements. The relations found by PIPS are given in the third column: they are obtained in less than a second on a typical PC, and are equivalent to those of the other tools.

Examples 1, 2 and 3 illustrate the need to compute disjunctive invariants. Examples 5 and 6 express linear invariants. Some periodic cases are presented in Section 3.2. Examples 4 and 7 and the Subway example [19] characterize automata with more complex invariants, and our results are equivalent to [10] (p. 115). However

our algorithm does not find accurate results with simple C encodings of automata such as the bakery mutual exclusion algorithm [6].

Using polyhedra instead of Presburger arithmetic, we do not claim to obtain more accurate results than others. Our philosophy is to use real-life cases, avoiding artificial or contrived examples. We only claim our simple and direct algorithm gets the same results as iterative approaches like widening.

The concept of abstract acceleration introduced by Gonnord in [10] is very similar in its goal to our algorithm, but it is implemented by pattern matching for different specific cases (see chapters 5 to 7 in [10]), whereas we can deal with function calls and any control construct as the loop body transformer is computed in a modular way. Also, the exploitation of the accelerated cycles is part of a heuristic and not a program transformation as in Section 3.5. And the final result is obtained iteratively. All examples found in [10] are successfully processed by our algorithm, including the swimming pool [9].

Kelly et al. [22] present an algorithm to compute the transitive closure of a relation encoded by a Presburger formulae. This heuristic includes the notion of *d-form relation* which leads to an explicit transitive closure. It is stated that any relation can be put in a *d-form* at the expense of accuracy. We show here that it is not necessary to put the relation into a *d-form* to obtain an explicit transitive closure. We explain how to transform any relation into constraints about the state evolution and finally we explain how to get precise results by postponing convex hull operations.

Bielecki et al. [2] describe a procedure to obtain exact non-linear transitive closures, but only for normalized relations written as systems of recurrence equations and solved using *Mathematica*. The related work is limited to [22] and a few examples are given. Regardless of the still unknown generality and practicality of this procedure, its results would require some processing to be used within an affine-base analyzer or algorithm as found in abstract interpretation and automatic parallelization.

Let us consider the code in Fig. 8 (left). Note that m is found monotonic in T^* in the second loop, thanks to the loop bound, and that the never ending loop is detected (the set of reaching states for `return` is empty), although closed forms are not computed for m . The same kind of results are obtained for a division in Fig. 8 (right). Variable m is found decreasing in the first loop and increasing in the second one, thanks to the loop bounds. Note also that m is not initialized, but that its final value is properly found in $[-1, 1]$.

In [27], Paige and Koenig propose finite differencing as a program optimization method that generalizes strength reduction. If the value of $f(x)$ is known and if $f(x+dx)$ is needed, is it possible to compute $f(x+dx) - f(x)$ faster than $f(x+dx)$? The simple case of strength reduction shows that our approaches are dual. We analyze the piece of code that computes $f(x+dx) - f(x)$ and we infer the function f . Paige and Koenig start with the code to evaluate $f(x)$ and infer the code to evaluate $f(x+dx) - f(x)$. Their technique was developed to optimize SETL code and to deal with functions over sets. The challenge for us would be to extend our

```

// T() empty
void multiply01() {
  // T(m) {m==1}
  int m = 1;
  // T(m) {m#init==1, m<=10}
  // P(m) {m==1}
  while (m<=10)
    // T(m) {m==2m#init, m#init<=10}
    // P(m) {m<=10}
    m = 2*m;
  // T(m) {11<=m, m#init<=m,
  // 11<=m#init, m#init<=20}
  // P(m) {11<=m, m<=20}
  while (m>=1)
    // T(m) {m==2m#init, 1<=m#init}
    // P(m) {11<=m}
    m = 2*m;
  // T() empty
  // P() empty
  return;
}

// T() {}
void divide01() {
  // T(m) {}
  int m;
  // T(m) {2<=m, m<=m#init, 2<=m#init}
  while (m>1) {
    // T(m) {m#init<=2m+1, 2m<=m#init, 2<=m#init}
    // P(m) {2<=m}
    m = m/2;
    // T(m) {1<=m}, P(m) {1<=m}
    printf("m=%d\n",m);
  }
  // T(m) {m+2<=0, m#init<=m, m#init+2<=0}
  // P(m) {m<=1}
  while (m<-1) {
    // T(m) {m#init<=2m+1, 2m<=m#init, m#init+2<=0},
    // P(m) {m+2<=0}
    m = m/2;
    // T(m) {m+1<=0}, P(m) {m+1<=0}
    printf("m=%d\n",m);
  }
  // T() {0<=m+1, m<=1}, P(m) {0<=m+1, m<=1}
  return;
}

```

Fig. 8. Beyond counters: multiply (left) & divide (right)

technique in a dual way to obtain predicates over arrays such as those found in [18].

Monotonicity analysis [29] has also been used to extend induction variable detection, the inverse transformation of strength reduction. Basically, assignment statements nested in loops are monotonic if the value assigned increases from one iteration to the next. The exact value of the difference is abstracted by its sign. This information does not lead to loop invariants but is useful for dependence testing [30] and for array bound checking. We could derive the same kind of information from T^+ by introducing the difference variables and by eliminating the program variables. The transformer we find for the contrived loop in Figure 7 of [29] is:

$$T(i,j,k,l,m,n,x,y,z) \{ \begin{array}{l} i==i_0+1, i==j-2, 2i==k-2, i+m\#init==l-2, l==m-1, \\ l==n-6, x==2y_0, 4y_0==z, 2<i, y<=2i+2, i<=3y+1 \end{array} \}$$

The first equation shows that i is increasing, hence j and k from the next two equations. Equations (4) and (5) lead to $m-m\#init=i+3$. The loop body precondition includes $i \geq 1$, so m is increasing, thus l and n as well. We also have $i-1 \leq 3y \leq 6i+6$. The lower and upper bounds are increasing but we cannot derive monotonicity information about y , nor about x and z . As the loop body is summarized, the monotonicity information is linked to the variables and not to the assignments, but this is exactly the same information.

5 Conclusion

A simple algorithm to compute affine invariants over integer scalar variables in while loops is presented. Its development and refinements have been mostly application driven, targeting the automatic program analysis and transformation domain. Its low complexity is key to addressing large scientific codes of up to 100 KLOC. Our experience shows that it performs well on standard program test cases, but not on complex automata whose states and transitions cannot be rewritten with simple C encodings.

This algorithm is more effective in finding loop invariants when inaccurate operations such as convex hulls and transitive closures are postponed as much as possible. If the analysis were exact, formulae such as $P^* = T^*(P^0)$ and $P^* = P^0 \sqcup T^+(P^0)$ would be equivalent. However, the formulae dealing with approximate transformers and preconditions are not and the best one must be chosen or a trade-off be made between accuracy and computational complexity. Developed formulae such as Eq.(9) correspond to peeling and unrolling the while loop in the computation. It is not compatible with Bourdoncle's heuristic, which aims at minimizing the number of widenings: better results are obtained by increasing the number of cycles and of transitive closures in order to delay convex hulls, as the closures are quite accurate.

When the program behavior is not affine, its affine approximations can be refined iteratively using the previous preconditions to obtain more accurate loop body transformers. This does not yield an algorithm as the iterations do not converge when the domain is not bounded (Section 3.4).

Our current implementation in PIPS is not fully satisfying as some extensions described in this paper are not available yet. They are not required often enough to justify the potential average slowdown and implementation time. Using a domain product instead of a unique general abstract domain could be investigated. Another idea would be to combine a widening heuristic and a derivative transitive closure algorithm. The later could be used for abstract acceleration in a widening heuristic along the lines of [10], but a combined approach is still in want of motivating test cases.

Acknowledgement

N. Halbwachs suggested years ago that our transitive closure algorithm should be published, however simple it was. It benefited from observations by B. Creusillet and N. Nguyen who performed tedious experiments and screened the results for missing information. P. Jouvelot and reviewers suggested many improvements in the presentation and in the related work. L. Gonnord helped us understand how her ASPIC tool works. P. Jouvelot formally proved the correctness of the code transformation of Section 3.5, see [1].

References

- [1] Ancourt, C., F. Coelho and F. Irigoin, *A modular static analysis approach to affine loop invariants detection (extended version)*, Technical Report A-419, CRI, MINES ParisTech (2010), parts to be published in NSAD 2010.
- [2] Bielecki, W., T. Klimek and K. Trifunovic, *Calculating exact transitive closure for a normalized affine integer tuple relation*, *Electronic Notes in Discrete Mathematics* **33** (2009), pp. 7 – 14, International Conference on Graph Theory and its Applications.
- [3] Boigelot, B., L. Bronne and S. Rassart, *Symbolic verification with periodic sets*, in: *Proc. 9th International Conference on Computer-Aided Verification, volume 1254, Lecture Notes in Computer Science* (1997), pp. 167–177.
- [4] Boigelot, B. and P. Wolper, *Symbolic verification with periodic sets*, in: *6th International Conference on Computer Aided Verification, number 808 in LNCS* (1994), pp. 55–67.

- [5] Bourdoncle, F., “Sémantiques des langages d’ordre supérieur et interprétation abstraite,” Ph.D. thesis, École polytechnique (1992).
- [6] Bultan, T., R. Gerber and W. Pugh, *Symbolic model checking of infinite state systems using presburger arithmetic*, in: *Proc. 9th International Conference on Computer-Aided Verification, volume 1254, Lecture Notes in Computer Science* (1997), pp. 400–411.
- [7] Comon, H. and Y. Jurski, *Multiple counters automata, safety analysis and presburger arithmetic*, in: *CAV ’98: Proceedings of the 10th International Conference on Computer Aided Verification* (1998), pp. 268–279.
- [8] Cousot, P. and N. Halbwachs, *Automatic discovery of linear restraints among variables of a program*, in: *POPL* (1978), pp. 84–96.
- [9] Fribourg, L. and H. Olsen, *Proving safety properties of infinite state systems by compilation into presburger arithmetic*, in: *CONCUR’97, LNCS 1243* (1997), pp. 213–227.
- [10] Gonnord, L., “Acceleration abstraite pour l’amélioration de la précision en analyse des relations linéaires,” Ph.D. thesis, Université Joseph Fourier, Grenoble, France (2007).
- [11] Gopan, D. and T. W. Reps, *Guided static analysis*, in: *Static Analysis, 14th International Symposium, SAS 2007*, 2007, pp. 349–365.
- [12] Gulavani, B. S., T. A. Henzinger, Y. Kannan, A. V. Nori and S. K. Rajamani, *Synergy: a new algorithm for property checking*, in: *SIGSOFT ’06*, 2006, pp. 117–127.
- [13] Gulwani, S., S. Jain and E. Koskinen, *Control-flow refinement and progress invariants for bound analysis*, in: *PLDI ’09*, 2009, pp. 375–385.
- [14] Gulwani, S. and N. Jojic, *Program verification as probabilistic inference*, *SIGPLAN Not.* **42** (2007), pp. 277–289.
- [15] Haghghat, M., “Symbolic analysis for parallelizing compilers,” Boston Kluwer Academic, 1996.
- [16] Halbwachs, N., “Détermination automatique de relations linéaires vérifiées par les variables d’un programme,” Ph.D. thesis, Université Scientifique et Médicale de Grenoble (1979).
- [17] Halbwachs, N., *Delay analysis in synchronous programs*, in: *Fifth Conference on Computer Aided Verification* (1993), pp. 4–13.
- [18] Halbwachs, N. and M. Péron, *Discovering properties about arrays in simple programs*, in: *PLDI ’08*, 2008, pp. 339–348.
- [19] Halbwachs, N., Y.-E. Proy and P. Roumanoff, *Verification of real-time systems using linear relation analysis*, *Form. Methods Syst. Des.* **11** (1997), pp. 157–185.
- [20] Irigoin, F., *Interprocedural analyses for programming environments*, in: *Environments and Tools for Parallel Scientific Computing* (1993), pp. 333–350.
- [21] Irigoin, F., P. Jouvelot and R. Triolet, *Semantical interprocedural parallelization: an overview of the pips project*, in: *ICS*, 1991, pp. 144–151.
- [22] Kelly, W., W. Pugh, E. Rosser and T. Shpeisman, *Transitive closure of infinite graphs and its applications*, *Int. J. Parallel Program.* **24** (1996), pp. 579–598.
- [23] Merchat, D., “Réduction du nombre de variables en analyse de relations linéaires,” Ph.D. thesis, Université Joseph Fourier (2005).
- [24] MINES-ParisTech, *PIPS*, <http://pips4u.org> (1989–2009), open source, under GPLv3.
- [25] Nguyen, N., “Efficient and Effective Software Verification for Scientific Applications Using Static Analysis and Code Instrumentation,” Ph.D. thesis, École des mines de Paris (2002).
- [26] Nguyen, T. V. N. and F. Irigoin, *Efficient and effective array bound checking*, *TOPLAS* **27** (2005), pp. 527–570.
- [27] Paige, R. and S. Koenig, *Finite differencing of computable expressions*, *TOPLAS* **4** (1982), pp. 402–454.
- [28] Pop, S., A. Cohen and G.-A. Silber, *Induction variable analysis with delayed abstractions*, Technical report, Ecole des mines de Paris, CRI A/367 (2005).
- [29] Spezialetti, M. and R. Gupta, *Loop monotonic statement*, *IEEE/TOSE* **21** (1995), pp. 497–505.
- [30] Wu, P., A. Cohen, J. Hoeflinger and D. Padua, *Monotonic evolution: an alternative to induction variable substitution for dependence analysis*, in: *ICS’01* (2001), pp. 78–91.
- [31] Zhou, C., C. A. R. Hoare and A. P. Ravn, *A calculus of durations*, *Information Processing Letters* (1991).

Appendix: Proof of Theorem 3.1

Lemma *If the elimination of variable z in a linear constraint system S by Fourier-Motzkin elimination produces the new system S' , the elimination of z in system S modified by multiplying all z coefficients by a positive integer produces the same S' .*

In other words, the Fourier-Motzkin elimination of a variable z is not perturbed if all coefficients of z are multiplied by the same positive constant.

Proof. Let S be $Ax \leq zb$ (since S is linear, there are no constant terms). S is decomposed into: $\{A^+x \leq zb^+, A^-x \leq -zb^-\}$ with $b = b^+ - b^-$.

The new constraints are built as $\{a_i^+x \leq zb_i^+, a_j^-x \leq -zb_j^-\}$, which leads to $b_j^-a^+x + b_i^+a_j^-x \leq 0$.

If all coefficients of z are multiplied by the same positive constant c , the decomposition in A^+ and A^- is not modified because c is positive and the new relations are equal to the old ones $\{a_i^+x \leq czb_i^+, a_j^-x \leq -czb_j^-\}$, and the inequation $cb_j^-a^+x + cb_i^+a_j^-x \leq 0$ which can be divided by c . \square

Proof of Theorem 3.1. $T(x, y)$ is a polyhedral transformer leading to $T'(y - x) = \{(y - x) | A(y - x) \leq b\}$ [no need to distinguish between equations and inequalities]. Then $(T^n)'$ verifies $(T^n)'(y - x) = \{(y - x) | A(y - x) \leq nb\}$. This is true for $n = 1$ and 2. If it is true for n , then $(T^{n+1})'(z - x) = (T^n)'(y - x) \wedge T'(z - y)$. Since $A(y - x) \leq nb$ and $A(z - y) \leq b$, $A(z - x) \leq (n + 1)b$. So the constraint systems defining T' and $(T^n)'$ differ only because the coefficients of z are multiplied by n . Using Lemma 1, the elimination of z (last step of the algorithm) leads to the same T^* . Hence all invariants of T^* are included in invariants of $(T^n)^*$. \square