On Intervals and Bounds in Bit-vector Arithmetic

Mikoláš Janota¹ and Christoph M. Wintersteiger¹

Microsoft Research {mikjan, cwinter}@microsoft.com

1 Introduction

Bit-vector arithmetic operations rely on modular arithmetic semantics, which often complicates algebraic manipulations. One such example are inequalities. Satisfiability of a system of modular inequalities is NP-complete [1]. In this memo we investigate a simple type of inequalities: only one variable is permitted and there are no multiplications. Satisfiability is solved in polynomial time by translating the problem into conjunctions of positive and negative intervals. This lets us decide systems of such inequalities but also identify uniqueness of solution and remove redundancies. This approach can be used to preprocess bit-vector problems and might also be useful in connection with other approaches, such as abstract interpretation [3,2].

2 Problem Definition

For the purpose of the memo we assume a fixed bit-with $w \in \mathbb{N}^+$. A bit vector variable is seen as an integer variable with possible values $0..(2^w - 1)$. Addition $a + {}^w b$ is defined as $(a + b) \mod 2^w$. Two types of comparisons are defined: *unsigned* $a \leq^{u} b$ and *signed* $a \leq^{s} b$. The unsigned is standard integer comparison, i.e., $a \leq^{u} b \Leftrightarrow a \leq b$. The signed uses the two's complement representation for signed numbers. Hence, we have $a \leq^{s} b \Leftrightarrow s(a) \leq s(b)$, where $s(a) \triangleq a < 2^{w-1}$? $a : a - 2^{w}$.

Consider a conjunction $\bigwedge_{1..m} \neg (a_i \odot b_i) \land \bigwedge_{m+1..n} (a_i \odot b_i)$, where $a_i \odot b_i$ is one of the following forms with x a bit-vector variable and c_1, c_2 constants.

- 1. $c_1 + {}^w x \leq^{\mathrm{u}} c_2 + {}^w x$ 2. $c_1 \leq^{\mathrm{u}} c_2 + {}^w x$ 3. $c_1 + {}^w x \leq^{\mathrm{u}} c_2$ 4. $x \leq^{\mathrm{s}} c_1$
- 5. $c_1 \leq^{\mathrm{s}} x$

Problem: Decide whether the input has a solution and if it has, determine if it is unique.

3 Solution

We show that every constraint of the form above can be represented as a positive or negative interval constraint. We write [a;b] and $\neg[a;b]$ for the positive and negative interval ranging from a to b inclusively, respectively. Hence, [a;b] for $0 \le a, b < 2^w$ denotes the set of values from a to b and $\neg[a;b]$ the values v with $(0 \le v < a) \lor (b < v < 2^w)$. Observe that $[a;b] \cup \neg[a;b] = [0;2^w-1]$. For a > b, it holds that $[a;b] = \emptyset$ and $\neg[a;b] = [0;2^w-1]$.

Expression	Condition	Interval		
$c_1 + {}^w x \leq {}^u c_2 + {}^w x$	$c_1 \leq c_2$	$\sim [-c_2; -c_1-1]$		
$c_1 + {}^w x \le {}^u c_2 + {}^w x$	$c_1 > c_2$	$[-c_1; -c_2-1]$		
$c_1 \leq^{\mathbf{u}} c_2 +^w x$	$c_1 < c_2$	$\sim [-c_2; c_1 - c_2 - 1]$		
$c_1 \leq^{\mathrm{u}} c_2 +^{w} x$	$c_1 \ge c_2$	$[c_1-c_2; -c_2-1]$		
$c_1 + {}^w x \le^{\mathrm{u}} c_2$	$c_1 \leq c_2$	$\sim [c_2 - c_1 + 1; -c_1 - 1]$		
$c_1 + {}^w x \le^{\mathrm{u}} c_2$	$c_1 > c_2$	$[-c_1; -c_1+c_2]$		
$x \leq^{\mathrm{s}} c_1$	$c_1 < 2^{w-1}$	$\sim [c_1 + 1; 2^{w-1} - 1]$		
$x \leq^{\mathrm{s}} c_1$	$c_1 \ge 2^{w-1}$	$[2^{w-1};c_1]$		
$c_1 \leq^{\mathrm{s}} x$	$c_1 < 2^{w-1}$	$[c_1; 2^{w-1} - 1]$		
$c_1 \leq^{\mathrm{s}} x$	$c_1 \ge 2^{w-1}$	$\sim [2^{w-1}; c_1 - 1]$		

Table 1. Translation of constraints.

Table 1 presents the conversion from the considered expressions to intervals. The table uses the convention that a negative number -c is automatically translated to the modular representative $2^w - c$.

The equivalence of the conversions is shown by case-splitting on whether the expressions $x + {}^{w}c_{i}$ overflow or not. So for instance $c_{1} + {}^{w}x \leq {}^{u}c_{2}$ with $c_{1} \leq c_{2}$ holds in two scenarios: (1) $c_{1} + {}^{w}x$ does not overflow and then it must be that $x \leq c_{2} - c_{1}$; (2) $c_{1} + {}^{w}x$ overflows and $c_{1} + {}^{w}x \leq {}^{u}c_{2}$ holds automatically because $(c_{1} + x - 2^{w}) \leq c_{1} \leq c_{2}$. The complement of the two scenarios is therefore $[c_{2} - c_{1} + 1; -c_{1} - 1]$.

Satisfiability as well as unique value property can be decided in polynomial time. Algorithm 1 present an algorithm that determines the upper and lower bound of the space determined by a set of intervals I. All the positive intervals are intersected obtaining thus lower bounds l and h, respectively. Negative intervals are processed in the increasing order of their lower bound, where the variable p is used as a pointer to the position where we may potentially find a point not covered by a negative interval. If the set of intervals is unsatisfiable then the algorithm returns [l'; h'] with l' > h', if there's a single solution then l' = h', and l' < h' means there are multiple solutions.¹

¹ The algorithm is a generalization of a standard solution to "The Interval Point Cover Problem".

Algorithm 1: Compute bounds

input : Set of intervals *I* **1** $P \leftarrow \{[a; b] \mid [a; b] \in I\}$ // get positive intervals $2 \ l \leftarrow P = \emptyset?0 : \min\{a \mid [a;b] \in P\}$ // get positive lower bound **3** $h \leftarrow P = \emptyset ? 2^w - 1 : \max \{b \mid [a; b] \in P\}$ // get positive upper bound 4 $N \leftarrow \{ \backsim[a; b] \mid \backsim[a; b] \in I \}$ // get negative intervals 5 $N \leftarrow \text{sort } N$ by first element **6** $p, l', h' \leftarrow l, l, l-1$ 7 for $\neg [a;b] \in N \cup \neg [2^w;2^w]$ do if p > h then break // space exhausted 8 if b < p then continue // redundant interval 9 // satisfiable portion if p < a then 10 if h' > l' then $l' \leftarrow p$ // first satisfiable point 11 $h' \leftarrow a - 1$ // update upper bound 12 $p \leftarrow b + 1$ // move onto next portion 13 14 return [l'; h']

3.1 Implementation and Redundancies

Algorithm 1 can be implemented in a straightforward fashion. Positive intervals do not need to be explicitly stored. Rather, we maintain a lower and upper bound L and U, respectively. These are updated with every new positive interval. These bounds let us also simplify negative intervals. An interval $\sim [a;b]$ with $a \leq L$ is simplified to $[b+1;2^{w-1}]$. Analogously, $\sim [a;b]$ with $b \geq u$ is simplified to [0;a-1].

Maintaining the upper and lower bounds lets us quickly detect if a new interval is redundant. A positive interval [a; b] is redundant if $L \leq a \wedge b \leq U$ and a negative interval $\sim [a; b]$ is redundant if $a > U \lor b < L$. Note that such detection of redundancies depends on the order of processing the intervals.

4 Examples

To measure the effect of the individual simplifications we construct several examples. (1) For $c_1 < c_2$, the conjunct $x + {}^w c_1 \leq x + {}^w c_2 \wedge x + {}^w c_2 \leq x + {}^w c_1$ is unsatisfiable. (2) For $x + {}^w c_1 \leq x + {}^w c_2 \wedge x + {}^w d_1 \leq x + {}^w d_2$ the second conjunct is redundant if $-c_2 \leq -d_2 \leq -d_1 - 1 \leq -c_1 - 1$. (3) For $x + {}^w c_1 \leq x + {}^w c_2 \wedge x + {}^w d_1 \leq x + {}^w d_2 \wedge x + {}^w a \leq x + {}^w b$ we obtain $x = -c_1$ if $-c_2 \leq -a \leq -c_1 - 1 \leq -b - 1 \leq -d_1 - 1$ and $-d_2 = -c_1 + 1$.

We generated 100 random instances for 32-bit numbers for each example and applied the standard bit-blasting technique in Z3 on them. Table 2 overviews the number of conflicts in the SAT solver.

Examples (1) and (2) are completely solved by Algorithm 1. Hence, implementing the algorithm in a preprocessor saves us the number of conflicts presented. For case (2) we compare the reduced and original instance, i.e., the redundant version contains two conjuncts whereas the reduced version only one.

Table 2. Conflict count.

Example	Avg.	Med	Min	Max
(1) unsat	26	36	3	100
(2) redundant	31	25	7	89
(2) reduced	29	37	3	92
(3) unique	21	40	0	123

The comparison shows that the biggest saving are in the case of unsatisfiable constraints and constraints that lead to unique solution. In the case of redundant intervals, however, the results do not indicate improvement.

5 Conclusion and Future Work

This memo shows how a system of inequalities over one variable and no multiplication can be represented as a set of positive and negative intervals. Naturally, such intervals are closed under negation, admit a polynomial satisfiability check and check for solution uniqueness. Preliminary evaluation suggests that such analysis enables significant speedups. In the future we plan to integrate this analysis into Z3 and investigate how could improve bit-vector solving by approximating solutions as intervals.

References

- Bjørner, N., Blass, A., Gurevich, Y., Musuvathi, M.: Modular difference logic is hard. CoRR abs/0811.0987 (2008), http://arxiv.org/abs/0811.0987
- Elder, M., Lim, J., Sharma, T., Andersen, T., Reps, T.: Abstract domains of affine relations. ACM Trans. Program. Lang. Syst. 36(4), 11:1–11:73 (Oct 2014), http: //doi.acm.org/10.1145/2651361
- Sharma, T., Thakur, A., Reps, T.: An abstract domain for bit-vector inequalities. Tech. Rep. TR1789, University of Wisconsin-Madison (2013), http://digital. library.wisc.edu/1793/65366