Inferring Congruence Equations Using SAT

Andy King¹ and Harald Søndergaard²

¹ Portcullis Computer Security Limited, Pinner, HA5 2EX, UK*

² The University of Melbourne, Victoria 3010, Australia.

Abstract. This paper proposes a new approach for deriving invariants that are systems of congruence equations where the modulo is a power of 2. The technique is an amalgam of SAT-solving, where a propositional formula is used to encode the semantics of a basic block, and abstraction, where the solutions to the formula are systematically combined and summarised as a system of congruence equations. The resulting technique is more precise than existing congruence analyses since a single optimal transfer function is derived for a basic block as a whole.

1 Introduction

Applications in compilation, optimisation and verification have motivated analyses that infer linear equality relationships [7,9,14] or linear congruence relationships [1,6,15] that hold between the variables of a program. For each point in a program, the former analyses discover systems of affine constraints of the form $\sum_{i=1}^n c_i x_i = d$ where $c_1, \dots, c_n, d \in \mathbb{Z}$ and $x_1, \dots, x_n$ are the program variables. The latter infer systems of congruence constraints of the form $(\sum_{i=1}^n c_i x_i) \bmod m = d$ where $m \in \mathbb{Z}$ is some modulus. These analyses accurately trace relationships between variables when the assignments that arise in a program can be modeled with a linear transformation. But this precludes meaningful analysis of programs that use bitwise operators; whether written in Java, C, or assembly language. The extreme approach of treating all operands of such operators as sequences of named bits, to track all bit interrelations, does not appear attractive, owing to the large number of Boolean variables involved. However, we show that a mixture of congruence analysis and Boolean reasoning does appear to be both feasible and able to generate bit-level invariants of great precision.

We draw inspiration from the domain of congruent equations modulo 2^w [15] and the affinity between this domain and the finite-nature of the underlying computer arithmetic, to propose an extreme-precision analysis which produces tight invariants for programs with non-linear, including bitwise, operations. The idea is to express the relationship between the bits in input variables and the bits of output variables for each basic block. This technique is not new within itself [8] and programs are now routinely reduced to very large systems of propositional constraints in bounded model checking [3,20]. Our main novel contribution is in the use of a SAT solver to incrementally compute a summary (affine relaxation) of the output variables, given a summary for the input variables. This new

* Andy King is on secondment from the University of Kent, CT2 7NF, UK

approach is capable of discovering invariants even for programs that apply bit-twiddling; programs that have thus far thwarted automatic analysis. Summaries that are systems of congruence equations modulo 2^w naturally fit into this mix of model checking and abstract interpretation because (technically) their ascending chain property constrains the number of times the SAT-solver can be reapplied and (philosophically) the propositional encoding also makes assumptions about the finite, modulo-nature of computer arithmetic. As in conventional abstract interpretation, the summaries enable all paths through the program to be considered systematically, without enforcing a bound on depth to which loops are explored. The approach to analysis is attractive because, quite apart from providing a bridging result between SAT solving and analysis, it can compute an optimal transfer function for a whole basic block, even when the block contains non-linear assignments. This is the key to the improved precision.

The paper is structured as follows: Section 2 illustrates the key ideas of the analysis using a worked example, demonstrating how a SAT-solver can be interleaved with a relaxation technique to compute a summary. This is the primary contribution of the paper. Section 3 shows how the lattice-theoretic join of two congruence systems can be summarised merely by syntactically rearranging matrices and computing a triangular form. This is another contribution of the paper. Section 4 discusses the relationship with the wider literature and Section 5 concludes.

2 Outline of the Method

In 1960 Peter Wegner [19] reported a fast bit counting algorithm. Expressed in the language C, the method counts the number of 1-bits in a word x , leaving the result in a variable c , as follows: `y = x; for (c = 0; y; c++) y &= y - 1;` Since the method is rather devious (and the explanation pre-dates the invariant assertion principle), one may want to derive an invariant that aids understanding of the code. This is challenging because the bit-twiddling cannot be modeled with a conventional affine assignment [15], that is, y is not updated with a value that is a linear combination of the values of the program variables. Furthermore, modeling the update as an assignment to an arbitrary value (a so-called non-deterministic assignment [15]) is too crude to derive a useful loop invariant.

One might think that these problems are insurmountable but we show that an invariant can be derived by modeling non-linear assignments as exact operations on sequences of bits and then computing optimal congruence abstractions for a composition of bit operations. The last point warrants elaboration: numeric analyses are usually presented in terms of programs which have already been abstracted through the use of assignments that are either affine or non-deterministic. This is adequate when working at the granularity of whole numbers but the best congruent abstraction of a bit-level operation, let alone a composition of them, is somewhat less obvious. We systematise the computation of an optimal abstraction and integrate this into the analysis itself.

In the rest of this section we sketch the approach:

1. A local, bit-precise transfer function is established for each basic block.
2. These Boolean functions are then used to build a set of recursive dataflow equations, expressing the program's overall runtime behaviour.
3. In the context of a finite set of w -bit variables, a closed form of the dataflow equations can be derived using Kleene iteration. In practice, however, this iteration may need to be interleaved with steps to relax constraints, and we propose a suitable relaxation to congruence equations.

2.1 Representing Bit-Level Semantics without Abstraction

It is possible to express the semantics of the basic blocks of a program, even to the bit-level, using Boolean constraints that relate the bits of its inputs to those of its outputs. But the problem is how to do so, retain tractability, and derive *loop invariants*, that is, not just explore loops to a fixed depth [8].

Let us draw Wegner's code as a flow diagram:

The program's basic blocks are the initial code ' $y := x; c := 0$ ', the loop body ' $\text{assume } y \neq 0; y := y \& (y - 1); c := c + 1$ ' and the loop exit ' $\text{assume } y = 0$ '. The exact semantics of these blocks can be described relationally, as systems of propositional constraints. The idea is to represent the input and output relationships across a basic block using two systems of propositional variables $x_0, \dots, x_{w-1}$ and $x'_0, \dots, x'_{w-1}$ (abbreviated to $\mathbf{x}$ and $\mathbf{x}'$) that encode the input and output state of each integer variable x . We assume a twos complement integer representation and let w denote the number of bits that make up an integer. The constraints generated for the example are:

$$\begin{aligned}
 \llbracket y := x; c := 0 \rrbracket &= (\bigwedge_{i=0}^{w-1} y'_i \leftrightarrow x_i) \wedge (\bigwedge_{i=0}^{w-1} \neg c'_i) \wedge (\bigwedge_{i=0}^{w-1} x'_i \leftrightarrow x_i) \\
 \llbracket \text{assume } y \neq 0; y := y \& (y - 1); c := c + 1 \rrbracket &= (\bigvee_{i=0}^{w-1} y_i) \wedge (\bigwedge_{i=0}^{w-1} y'_i \leftrightarrow (y_i \wedge \bigvee_{j=0}^{i-1} y_j)) \\
 &\quad \wedge (\bigwedge_{i=0}^{w-1} c'_i \leftrightarrow (c_i \oplus \bigwedge_{j=0}^{i-1} c_j)) \wedge (\bigwedge_{i=0}^{w-1} x'_i \leftrightarrow x_i) \\
 \llbracket \text{assume } y = 0 \rrbracket &= (\bigwedge_{i=0}^{w-1} \neg y_i) \wedge (\bigwedge_{i=0}^{w-1} x'_i \leftrightarrow x_i) \wedge (\bigwedge_{i=0}^{w-1} y'_i \leftrightarrow y_i) \wedge (\bigwedge_{i=0}^{w-1} c'_i \leftrightarrow c_i)
 \end{aligned}$$

where $\oplus$ denotes exclusive-or. Elsewhere [11] we explain how these constraints can be generated automatically from the program. Suffice it to say that the constraint for an assignment $\mathbf{x} := \mathbf{y} + \mathbf{z}$ is derived by considering a cascade of full adders using intermediate carry bits $\mathbf{b}$,

$$\left(\bigwedge_{i=0}^{w-1} x'_i \leftrightarrow y_i \oplus z_i \oplus b_i \right) \wedge \neg b_0 \wedge \left(\bigwedge_{i=1}^{w-1} b_i \leftrightarrow (y_{i-1} \wedge z_{i-1}) \vee (y_{i-1} \wedge b_{i-1}) \vee (z_{i-1} \wedge b_{i-1}) \right)$$

together with constraints to express that variables other than $\mathbf{x}$ do not change. The variables $\mathbf{b}$ are existentially quantified, and the formula can be simplified using standard Boolean quantifier elimination.

Compound expressions are handled by introducing temporary variables $\mathbf{s}$ and $\mathbf{t}$ to hold intermediate results and then applying renaming. For example,

$$\llbracket y := y \ \& \ (y - 1) \rrbracket = \rho_{\mathbf{s}', \mathbf{t}}(\llbracket s := y - 1 \rrbracket \wedge \llbracket y := y \ \& \ t \rrbracket)$$

The renaming step $\rho_{\mathbf{s}', \mathbf{t}}$ replaces the output variables $\mathbf{s}'$ of the first statement with the input variables $\mathbf{t}$ of the second. Again, the compound formula can be simplified by eliminating the remaining intermediate variables $\mathbf{s}, \mathbf{t}, \mathbf{t}'$.

2.2 Setting Up the Dataflow Equations

The relational semantics for the basic blocks allows us to derive the states that are reachable at program points 0, 2 and 5. They are obtained as the least (strongest) solution to the following recursive equations:

$$\begin{aligned} f_0 &= 1 \\ f_2 &= \rho_{\mathbf{v}', \mathbf{v}}(\pi_{\mathbf{v}'}(f_0 \wedge \llbracket y := x; c := 0 \rrbracket)) \vee \\ &\quad \rho_{\mathbf{v}', \mathbf{v}}(\pi_{\mathbf{v}'}(f_2 \wedge \llbracket \text{assume } y \neq 0; y := y \& (y - 1); c := c + 1 \rrbracket)) \\ f_5 &= \rho_{\mathbf{v}', \mathbf{v}}(\pi_{\mathbf{v}'}(f_2 \wedge \llbracket \text{assume } y = 0 \rrbracket)) \end{aligned}$$

where $\mathbf{v}$ and $\mathbf{v}'$ are the input and output variables, that is, $\mathbf{v} = \mathbf{c} \cdot \mathbf{x} \cdot \mathbf{y}$ and $\mathbf{v}' = \mathbf{c}' \cdot \mathbf{x}' \cdot \mathbf{y}'$. The Boolean functions f_0 , f_2 and f_5 represent sets of reachable states, for example, a state $\sigma = \{c_0 \mapsto 0, \dots, c_{w-1} \mapsto 0, x_0 \mapsto 1, \dots, x_{w-1} \mapsto 0, y_0 \mapsto 1, \dots, y_{w-1} \mapsto 0\}$ is considered to be reachable at program point 2 iff σ satisfies f_2 . The projection operation $\pi_{\mathbf{v}'}(f) = f'$ computes the Boolean function f' by eliminating, by existential quantification, any propositional variable y from f that does not occur in the system $\mathbf{v}'$. For instance, $\pi_{\langle x_1, x_2 \rangle}(f) = \exists_y(f) = x_1 \leftrightarrow x_2$ when $f = (x_1 \leftrightarrow y) \wedge (y \leftrightarrow x_2)$. Projection is used to derive a function that only expresses relations between the output variables $\mathbf{v}'$. The renaming operation $\rho_{\mathbf{v}', \mathbf{v}}(f) = f'$ constructs a function f' by replacing each output variable y' in f with its counterpart input variable y , for example, $\rho_{\mathbf{v}', \mathbf{v}}(c'_0 \wedge (x'_1 \oplus y'_1)) = c_0 \wedge (x_1 \oplus y_1)$. Iteration can be used to compute f_2 from the predetermined $f_0 = 1$ and once f_2 is known, f_5 can be derived. For f_2 the iterates start:

$$\begin{aligned} f_2^0 &= 0 \\ f_2^1 &= f_2^0 \quad \vee \quad (\bigwedge_{i=0}^{w-1} x_i \leftrightarrow y_i) \wedge (\bigwedge_{i=0}^{w-1} \neg c_i) \\ f_2^2 &= f_2^1 \quad \vee \quad (\bigvee_{i=0}^{w-1} x_i) \wedge (\bigwedge_{i=0}^{w-1} y_i \leftrightarrow (x_i \wedge \bigvee_{j=0}^{i-1} x_j)) \wedge c_0 \wedge (\bigwedge_{i=1}^{w-1} \neg c_i) \\ f_2^3 &= f_2^2 \quad \vee \quad \dots \end{aligned}$$

This sequence will eventually stabilise because a bounded number (2^{3w}) of Boolean functions are definable over $\mathbf{v}$. However, the $\mathbf{c}$ variables will enumerate all 2^w bit patterns and therefore at least 2^w iterates will be computed. This will take an impractically long time, even for $w = 16$. There is also an issue of space. A Boolean function can be represented as an ROBDD but the size of

an ROBDD can be exponentially large in the number of variables (even with dynamic variable reordering [2]), and this is a pressing issue when w propositional variables are needed to represent each integer variable. Tractability can be recovered by approximating [17] or widening [10] an ROBDD when it becomes intolerably large. This would replace an ROBDD with one that could be stored more compactly and yet represented a larger set of states. The problem with this approach is that the ROBDD widenings that have been proposed thus far do not preserve sufficient information to infer useful loop invariants.

2.3 Abstracting Bit-Level Inputs and Outputs with Congruences

Considerations of tractability dictate that we look for principled ways of over-approximating solutions to systems of equations of the form

$$f = \bigvee_{m=1}^n \rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(f_m \wedge f'_m)) \quad (1)$$

without giving up too much bit-level information. We suggest that this can be achieved by restricting f , as well as each f_m , to a class of functions that can be expressed as conjunctions of congruence equations modulo m , where m is a power of 2. Each function f_m summarises the inputs to one of n basic blocks and the function f summarises all (the join of) the outputs of the n blocks. *No constraint is placed on the generality of any of the f'_m formulae.* This means that no abstraction needs to be applied to a function that describes the relational semantics of a basic block—this description is kept bit-precise.

How to solve systems of the form (1) under the restrictions just mentioned? The rest of this section explains the idea, based on the Wegner example. Let $x \equiv_{2^w} y$ abbreviate $x = y + k2^w$ for some integer multiplier k . Observe that each of the equations $t \equiv_{2^w} t' + t''$, $t \equiv_{2^w} ny$, $t \equiv_{2^w} n$ and the disequation $t \not\equiv_{2^w} n$ can be expressed as propositional constraints when the t variables are w -bit, y is 1-bit and n is an integer. This is a consequence of the modulus being a power of 2. For instance, $t \equiv_{2^w} ny$ and $t \not\equiv_{2^w} ny$ can be expressed as $\bigwedge_{i=0}^{w-1} t_i \leftrightarrow (n_i \wedge y)$ and $\bigvee_{i=0}^{w-1} t_i \oplus (n_i \wedge y)$ where $n \equiv_{2^w} \sum_{i=0}^{w-1} 2^i n_i$ and $t \equiv_{2^w} \sum_{i=0}^{w-1} 2^i t_i$. Moreover, an equation $\sum_{i=1}^k n_i y_i \equiv_{2^w} n$ can be reduced by $\sum_{i=1}^j n_i y_i \equiv_{2^w} t$, $\sum_{i=j+1}^k n_i y_i \equiv_{2^w} t'$, $t + t' \equiv_{2^w} t''$ and $t'' \equiv_{2^w} n$ using fresh w -bit variables t , t' , and t'' , and hence also reduced to a propositional system. Any disequation $\sum_{i=1}^k m_i y_i \not\equiv_{2^w} n$ can similarly be described propositionally. Henceforth let $\llbracket \sum_{i=1}^k n_i y_i \equiv_{2^w} n \rrbracket$ and $\llbracket \sum_{i=1}^k n_i y_i \not\equiv_{2^w} n \rrbracket$ denote such propositional encodings.

To illustrate the value of these encodings, let $w = 8$ and consider computing $f_2^2 = f_2^1 \vee \rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(f_2^1 \wedge g))$ where $f_2^1 = \rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(1 \wedge \llbracket y := x; c := 0 \rrbracket))$ and $g = \llbracket \text{assume } y \neq 0; y := y \& (y - 1); c := c + 1 \rrbracket$. The encodings are used to direct the generation of truth assignments for the function $f_2^1 \wedge g$. Truth assignments are generated so as to incrementally derive the most precise congruence system describing both f_2^1 and $\rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(f_2^1 \wedge g))$. This system is used as the iterate, rather than the function f_2^1 itself. Since the function $f_2^0 = 0$ can be represented

as a congruence system, namely $0 \equiv_{256} 1$, this construction ensures that all iterates are congruences. The number of iterates is bounded: the length of an increasing chain of congruences is at most 192, that is, the product of the power $w = 8$ and the maximum number, 24 ($3w$), of variables in each system [15].

The function f_2^1 falls into the class of formulae that can be represented congruently. This is because the satisfying assignments of f_2^1 coincide with those of the formula $(\bigwedge_{i=0}^7 \llbracket c_i \equiv_{256} 0 \rrbracket) \wedge (\bigwedge_{i=0}^7 \llbracket x_i \equiv_{256} y_i \rrbracket)$ on $\mathbf{c}$, $\mathbf{x}$ and $\mathbf{y}$. Hence computing $\rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(f_2^1 \wedge g))$ is equivalent to computing $\rho_{\mathbf{y}', \mathbf{y}}(\pi_{\mathbf{y}'}(g'))$ where $g' = (\bigwedge_{i=0}^7 \llbracket c_i \equiv_{256} 0 \rrbracket) \wedge (\bigwedge_{i=0}^7 \llbracket x_i \equiv_{256} y_i \rrbracket) \wedge g$. This is convenient, because it permits the encodings to be demonstrated on a representative, non-trivial example, namely the derivation of f_2^2 . To see how the congruence system for f_2^2 is incrementally constructed, observe that any model of the Boolean function

$$g' \wedge ((\bigvee_{i=0}^7 \llbracket c'_i \not\equiv_{256} 0 \rrbracket) \vee (\bigvee_{i=0}^7 \llbracket x'_i \not\equiv_{256} y'_i \rrbracket)) \quad (2)$$

defines a run of the block with an entry state that is described by f_2^1 but whose exit state is *not* characterised by f_2^1 . For example, the truth assignment

$$\left\{ \begin{array}{l} c_0 \mapsto 0, c_1 \mapsto 0, \dots, c_7 \mapsto 0, x_0 \mapsto 0, \dots, x_6 \mapsto 0, x_7 \mapsto 1, y_0 \mapsto 0, \dots, y_7 \mapsto 1, \\ c'_0 \mapsto 1, c'_1 \mapsto 0, \dots, c'_7 \mapsto 0, x'_0 \mapsto 0, \dots, x'_6 \mapsto 0, x'_7 \mapsto 1, y'_0 \mapsto 0, \dots, y'_7 \mapsto 0 \end{array} \right\}$$

satisfies (2) and demonstrates that when $\mathbf{c}$, $\mathbf{x}$, and $\mathbf{y}$ assume values of 0, 128 and 128 on entry to the block, they can take values of 1, 128 and 0 on exit from the block (assuming an unsigned representation). By construction, the output state is not summarised by f_2^1 and therefore f_2^1 needs to be enlarged to accommodate this state. Since the output state can be represented in congruence form as

$$c_0 \equiv_{256} 1 \wedge (\bigwedge_{i=1}^7 c_i \equiv_{256} 0) \wedge (\bigwedge_{i=0}^6 x_i \equiv_{256} 0) \wedge x_7 \equiv_{256} 1 \wedge (\bigwedge_{i=0}^7 y_i \equiv_{256} 0) \quad (3)$$

this system and that for f_2^1 can be joined to obtain the summary

$$(\bigwedge_{i=1}^7 c_i \equiv_{256} 0) \wedge (\bigwedge_{i=0}^6 x_i \equiv_{256} y_i) \wedge x_7 \equiv_{256} c_0 + y_7 \quad (4)$$

A model for the formula (2) can be found using standard techniques [16], translating the formula into an equi-satisfiable conjunctive normal form (CNF) representation, and presenting the CNF formula to any SAT-solver. The join can be computed by translating the two congruence systems to their sets of generators, taking the union of the two sets, then converting the union to a new congruence system [1,6,15]. Alternatively, the join can be obtained by relaxing a system of congruences constructed syntactically from the two input systems (see Section 3). Either way, whether the join describes all possible output states can be determined by solving the Boolean formula

$$g' \wedge ((\bigvee_{i=1}^7 \llbracket c_i \not\equiv_{256} 0 \rrbracket) \vee (\bigvee_{i=0}^6 \llbracket x_i \not\equiv_{256} y_i \rrbracket) \vee \llbracket x_7 \not\equiv_{256} c_0 + y_7 \rrbracket) \quad (5)$$

This formula is satisfied, for example, by a truth assignment $\{\dots, c'_0 \mapsto 1, c'_1 \mapsto 0, \dots, c'_7 \mapsto 0, x'_0 \mapsto 0, \dots, x'_5 \mapsto 0, x'_6 \mapsto 1, x'_7 \mapsto 0, y'_0 \mapsto 0, \dots, y'_7 \mapsto 0\}$ from which the following congruence system can be derived:

$$c_0 \equiv_{256} 1 \wedge (\bigwedge_{i=1}^7 c_i \equiv_{256} 0) \wedge (\bigwedge_{i=0, i \neq 6}^7 x_i \equiv_{256} 0) \wedge x_6 \equiv_{256} 1 \wedge (\bigwedge_{i=0}^7 y_i \equiv_{256} 0) \quad (6)$$

Note that the assignments to the input variables, as well as any temporary variables introduced in CNF conversion [16], are inconsequential for constructing the congruence system. Disregarding these assignments amounts to projecting onto the output variables. Notice too that the system is expressed in terms of unprimed variables, even though it encodes an output state. Constructing the congruence system thus involves renaming as well as projection, though both operations are performed on truth assignments, at which level they collapse to computationally trivial operations. Joining the previous summary (4) with the new system (6) gives the new summary

$$(\bigwedge_{i=1}^7 c_i \equiv_{256} 0) \wedge (\bigwedge_{i=1}^5 x_i \equiv_{256} y_i) \wedge x_6 + x_7 \equiv_{256} c_0 + y_6 + y_7 \quad (7)$$

Continuing this way, we obtain a sequence of Boolean formulae $h_0, h_1, h_2, \dots$, the first two of which are (2) and (5), and where, more generally, h_j is

$$g' \wedge ((\bigvee_{i=1}^7 \llbracket c_i \not\equiv_{256} 0 \rrbracket) \vee (\bigvee_{i=0}^{7-j} \llbracket x_i \not\equiv_{256} y_i \rrbracket) \vee \llbracket \sum_{i=7-j+1}^7 x_i \not\equiv_{256} c_0 + \sum_{i=7-j+1}^7 y_i \rrbracket)$$

Of these, $h_1, \dots, h_6$ are satisfiable, but h_7 is not, so the system

$$(\bigwedge_{i=1}^7 c_i \equiv_{256} 0) \wedge \sum_{i=0}^7 x_i \equiv_{256} c_0 + \sum_{i=0}^7 y_i \quad (8)$$

summarises all reachable output states when the input states are described by f_2^1 . The next iterate f_2^2 is then assigned to this system which is the most precise congruence system that describes the set of output states given an input state drawn from f_2^1 .

The method is not sensitive to how the relational semantics is presented. This contrasts with previous analyses which critically depend on how the statements of a program are translated into, say, affine assignments. This is particularly pertinent when deriving invariants for assembler or obfuscated code [18].

Of course, thus far, only f_2^2 has been derived. By repeating the above process with an updated input state we obtain the sequence of iterates:

$$\begin{aligned} f_2^3 &= (\bigwedge_{i=2}^7 c_i \equiv_{256} 0) \wedge \sum_{i=0}^7 x_i \equiv_{256} c_0 + 2c_1 + \sum_{i=0}^7 y_i \\ f_2^4 &= (\bigwedge_{i=3}^7 c_i \equiv_{256} 0) \wedge \sum_{i=0}^7 x_i \equiv_{256} c_0 + 2c_1 + 4c_2 + \sum_{i=0}^7 y_i \\ f_2^6 &= f_2^5 = (\bigwedge_{i=4}^7 c_i \equiv_{256} 0) \wedge \sum_{i=0}^7 x_i \equiv_{256} c_0 + 2c_1 + 4c_2 + 8c_3 + \sum_{i=0}^7 y_i \end{aligned}$$

Interestingly, although the derivation of f_2^2 requires 8 calls to a SAT-solver and 7 join computations, the iterates f_2^3 , f_2 and f_2^5 each require just two calls to a solver and one join. To check stability, that is, deduce $f_2^6 = f_2^5$, requires one call to a solver, hence 15 invocations are required in total, the largest of which involves 4507 variables and 11648 clauses (though more compact CNF conversion is possible [16]). Nevertheless, the longest time that it takes to solve any instance is 0.61 ms (wall-time) using SAT4J version 1.5 [12] on a 2.4 GHz MacBook Pro. Even without deriving $f_5 = \sum_{i=0}^7 x_i \equiv_{256} c_0 + 2c_1 + 4c_2 + 8c_3 \wedge (\bigwedge_{i=4}^7 c_i \equiv_{256} 0)$, it is now evident that Wegner's bit-twiddling algorithm assigns to the variable $\mathbf{c}$ the number of bits which are set in the variable $\mathbf{x}$. As far as we aware, no other analysis is capable of deriving such an invariant. Note that the invariant includes coefficients of 4 and 8, and thus precision would be degraded if a modulo of 2 was employed rather than the word-level modulo of 2^w .

3 Joining Congruence Equations

Section 2.1 outlined the translation of basic blocks into Boolean formulae. That component preserves all information within a basic block, which is key to reasoning about non-linear operations such as bit-twiddling. We now describe the complementary component which produces the join of two congruence systems. It discards information, which is key to retaining tractability. The join ensures that the summaries reside in a finite ascending chain so they cannot be weakened forever; the maximal chain length is w^2n [15], as wn (propositional) variables are needed to represent the state of n (integer) variables of width w .

Recent work has exploited how congruence systems can be represented by sets of generators that span the solution space of the congruence system [15]. This representation is useful because it reduces the join operation to set union. However, our refined form of analysis relies on a translation mechanism from an equation $\sum_{i=1}^k n_i y_i \equiv_{2^w} n$ to a formula $\llbracket \sum_{i=1}^k n_i y_i \equiv_{2^w} n \rrbracket$ that becomes more convoluted when the generator representation is adopted. Thus it is convenient to compute the join whilst representing the input and output systems as a conjunction of equations. This can be achieved by reformulating the join of two systems as a projection operation which can, in turn, be computed by calculating a triangular form. This section explains these steps.

Basics. To state the algorithmic results of this section, it is necessary to recall some mathematical concepts and notation. The set of congruence classes modulo m is defined $\mathbb{Z}_m = \{[n] \mid n \in \mathbb{Z}\}$ where $[n] = \{n' \in \mathbb{Z} \mid n \equiv_m n'\}$ and $\equiv_m$ denotes equivalence modulo m . Henceforth, we blur the distinction between a class $[n]$ and its representative element n . The 2-fold Cartesian product $\mathbb{Z}_m^2$ is defined $\mathbb{Z}_m^2 = \mathbb{Z}_m \times \mathbb{Z}_m$ and the k -fold product $\mathbb{Z}_m^k$ is likewise defined. If $S_1, S_2 \subseteq \mathbb{Z}_m^k$ then their (Minkowski) sum is $S_1 + S_2 = \{\mathbf{x} \in \mathbb{Z}_m^k \mid \mathbf{x}_i \in S_i \wedge \mathbf{x} \equiv_m \mathbf{x}_1 + \mathbf{x}_2\}$. If $\lambda \in \mathbb{Z}$ and $S \subseteq \mathbb{Z}_m^k$, then $\lambda S = \{\mathbf{x} \in \mathbb{Z}_m^k \mid \mathbf{x}' \in S \wedge \mathbf{x} \equiv_m \lambda \mathbf{x}'\}$. Moreover, the linear closure of S is $\text{linear}(S) = \{\sum_{i=1}^{\ell} \lambda_i \mathbf{x}_i \mid \lambda_i \in \mathbb{Z} \wedge \mathbf{x}_1, \dots, \mathbf{x}_{\ell} \in S\}$.


```

1: procedure triangular(in:  $S$ , out:  $t$ )
2:    $t := \lambda \ell. \perp$ 
3:   let  $\{a_1, \dots, a_s\} = S$ 
4:   for  $i := 1$  to  $s$  do
5:      $\ell := \text{leading}(a_i)$ 
6:     while  $(\ell > 0 \wedge \ell \in \text{dom}(t))$ 
7:        $a' := t(\ell)$ 
8:        $p := \text{power}(\pi_\ell(a_i))$ 
9:        $p' := \text{power}(\pi_\ell(a'))$ 
10:      if  $p \geq p'$  then
11:         $a_i := (\pi_\ell(a')/2^{p'})a_i - 2^{p-p'}a'$ 
12:      else
13:         $t := t[\ell \mapsto a_i]$ 
14:         $a_i := (\pi_\ell(a_i)/2^p)a' - 2^{p-p}a_i$ 
15:      endif
16:       $\ell := \text{leading}(a_i)$ 
17:    endwhile
18:    if  $\ell > 0$  then  $t := t[\ell \mapsto a_i]$ 
19:  endfor
20: endprocedure
    
```

Fig. 1. The triangularisation method of Müller-Olm and Seidl [15]

The modules of $\mathbb{Z}_m^k$ are those subsets of $\mathbb{Z}_m^k$ that are closed under linear combination, that is, $\text{Module}_m^k = \{M \subseteq \mathbb{Z}_m^k \mid \text{linear}(M) = M\}$. The affine subsets of $\mathbb{Z}_m^k$ are translated modules, that is, $\text{Affine}_m^k = \{\{x\} + M \mid x \in \mathbb{Z}_m^k \wedge M \in \text{Module}_m^k\}$. The affine closure of S is the smallest affine space that encloses S and is thus defined $\text{affine}(S) = \bigcap \{S' \in \text{Affine}_m^k \mid S \subseteq S'\}$.

Example 1. Observe that $\emptyset$ and $\{0\}$ are closed under linear combination, whence $\emptyset \in \text{Module}_m^k$ and $\{0\} \in \text{Module}_m^k$. As $\emptyset \in \text{Module}_m^k$, we have $\emptyset \in \text{Affine}_m^k$. Moreover, since $\{0\} \in \text{Module}_m^k$, it follows that $\{x\} \in \text{Affine}_m^k$ for all $x \in \mathbb{Z}_m^k$.

Triangularisation. While congruence equations appear as good companions for a bit-level relational semantics, Gaussian elimination cannot be immediately applied to compute a triangular form for such equations because of the need to deal with zero divisors [15]. Müller-Olm and Seidl have thus devised a triangularisation algorithm that, given an input system $Ax \equiv_{2^w} b$, computes output system $A'x \equiv_{2^w} b'$ where $A' = [a_{i,j}]$ and $a_{i,j} = 0$ whenever $i > j$. Figure 1 gives the algorithm and Example 2 illustrates what the algorithm will compute for an example. (This example will, in turn, support Example 4 which demonstrates how join can be straightforwardly realised.) In the description of the algorithm, $\text{leading}(a)$ returns -1 if $a = 0$ and otherwise the position of the first non-zero element of the vector a ; $\pi_\ell(a)$ extracts the ℓ 'th element from a ; and $\text{power}(n)$ returns the largest integer p such that 2^p divides n .

Example 2. The input and output to triangularisation procedure are $A\mathbf{x} \equiv_2 \mathbf{b}$ and $A'\mathbf{x} \equiv_2 \mathbf{b}'$ respectively:

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} \quad \mathbf{b} = \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \quad A' = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \quad \mathbf{b}' = \begin{bmatrix} 1 \\ 1 \\ 0 \\ 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 1 \end{bmatrix}$$

By reasoning about upper triangular form, one can argue that any subset of $\mathbb{Z}_m^k$ that is closed under affine combination, can be represented congruently:

Proposition 1. $S \in \text{Affine}_m^k$ iff there exists a congruence system $A\mathbf{x} \equiv_m \mathbf{b}$ such that $S = \{\mathbf{x} \in \mathbb{Z}_m^k \mid A\mathbf{x} \equiv_m \mathbf{b}\}$.

Projection. Quite apart from establishing this result, upper triangular form provides a way of computing arbitrary projections. Projection onto the i 'th element of a k -ary vector is defined $\pi_i(\langle x_1, \dots, x_k \rangle) = x_i$. Single element projections can be composed so that if $1 \leq i_1 < \dots < i_j \leq k$ then the j -ary vector $\langle \pi_{i_1}(\mathbf{x}), \dots, \pi_{i_j}(\mathbf{x}) \rangle$ is also a projection in that it also discards information pertaining to certain dimensions. Projection of an affine space is also affine:

Proposition 2. Let $S \in \text{Affine}_m^k$ and $1 \leq i_1 < \dots < i_j \leq k$. Then $T \in \text{Affine}_m^j$ where $T = \{\langle \pi_{i_1}(\mathbf{x}), \dots, \pi_{i_j}(\mathbf{x}) \rangle \mid \mathbf{x} \in S\}$.

If $A = [a_{i,j}]$ is in upper triangular form, the projection of $A\mathbf{x} \equiv_m \mathbf{b}$ onto a suffix $\mathbf{y} = \langle x_i, \dots, x_k \rangle$ of $\mathbf{x}$ is found very easily. Suppose row j is the top-most row of A in which $\langle a_{j,1}, \dots, a_{j,i-1} \rangle = \mathbf{0}$. Then the projection onto $\mathbf{y}$ is

$$\begin{bmatrix} a_{j,i} & \cdots & a_{j,k} \\ \vdots & & \vdots \\ a_{s,i} & \cdots & a_{s,k} \end{bmatrix} \mathbf{y} \equiv_m \begin{bmatrix} \pi_j(\mathbf{b}) \\ \vdots \\ \pi_s(\mathbf{b}) \end{bmatrix}$$

Example 3. Projecting $A\mathbf{x} \equiv_2 \mathbf{b}$ of Example 2, or equivalently the system $A(t)\mathbf{x} \equiv_2 \mathbf{b}(t)$, onto $\mathbf{y}_i = \langle x_i, \dots, x_{11} \rangle$ for $i = 7, 9$ and 10 yields:

$$\begin{bmatrix} 1 & 0 & 0 & 0 & 0 \\ 0 & 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 & 0 \end{bmatrix} \mathbf{y}_7 \equiv_2 \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} \quad [1 \ 1 \ 0] \mathbf{y}_9 \equiv_2 [1] \quad [0 \ 0] \mathbf{y}_{10} \equiv_2 [0] \quad (\text{or } 1)$$

Given a congruence system $A\mathbf{x} \equiv_m \mathbf{b}$, it is possible to project onto any subset of $\mathbf{x}$ merely by reordering the rows of A in synchronicity with the elements of $\mathbf{b}$, prior to computing the triangular form.

Join. We finally show how the join can be reduced to computing a projection (a relaxation) which, in turn, amounts to deriving an upper triangular form.

Proposition 3. Let $S_i = \{\mathbf{x} \in \mathbb{Z}_m^k \mid A_i \mathbf{x} \equiv_m \mathbf{b}_i\}$ and

$$A = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 \\ -\mathbf{b}_1 & 0 & A_1 & 0 & 0 \\ 0 & -\mathbf{b}_2 & 0 & A_2 & 0 \\ 0 & 0 & -I & -I & I \end{bmatrix} \quad S = \left\{ \mathbf{x} \in \mathbb{Z}_m^k \mid \exists \sigma_i \in \mathbb{Z}_m. \exists \mathbf{x}_i \in \mathbb{Z}_m^k. A \begin{bmatrix} \sigma_1 \\ \sigma_2 \\ \mathbf{x}_1 \\ \mathbf{x}_2 \\ \mathbf{x} \end{bmatrix} \equiv_m \begin{bmatrix} 1 \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \\ \mathbf{0} \end{bmatrix} \right\}$$

Then $S = \text{affine}(S_1 \cup S_2)$ if $S_1 \neq \emptyset$ and $S_2 \neq \emptyset$.

If a system $A_i \mathbf{x} \equiv_{2^w} \mathbf{b}_i$ has a solution set S_i , then the join of $A_1 \mathbf{x} \equiv_{2^w} \mathbf{b}_1$ and $A_2 \mathbf{x} \equiv_{2^w} \mathbf{b}_2$ is a system whose solutions coincide with $\text{affine}(S_1 \cup S_2)$. Proposition 3 states that such a system can be obtained by rearranging A_1 and A_2 to form a new matrix A and then eliminating variables.

Example 4. Consider the join of $A_1 \mathbf{x} \equiv_2 \mathbf{b}_1$ and $A_2 \mathbf{x} \equiv_2 \mathbf{b}_2$ where $\mathbf{x} = \langle x, y, z \rangle$

$$A_1 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \end{bmatrix} \quad \mathbf{b}_1 = \begin{bmatrix} 0 \\ 1 \end{bmatrix} \quad A_2 = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad \mathbf{b}_2 = \begin{bmatrix} 1 \\ 0 \\ 1 \end{bmatrix}$$

As well as minimising the size of coefficients and thereby making the presentation of large matrices manageable, a by-product of $2^w = 2$ is that $-I \equiv_2 I$ and $\mathbf{b}_i \equiv_2 -\mathbf{b}_i$. Using this, the combined system of Proposition 3 is formed—it is given below, on the left. On the right is the triangular system derived in Example 2, and we conclude that the join is $x + y \equiv_2 1$.

$$\begin{array}{c|c|c|c|c|c} \hline 1 & 1 & 0 & 0 & 0 & 0 \\ \hline 0 & 0 & 1 & 0 & 0 & 0 \\ \hline 1 & 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 1 & 0 & 0 & 0 & 1 \\ \hline 0 & 0 & 1 & 0 & 0 & 1 \\ \hline 0 & 0 & 0 & 1 & 0 & 0 \\ \hline 0 & 0 & 0 & 1 & 0 & 1 \\ \hline 0 & 0 & 0 & 0 & 1 & 0 \\ \hline 0 & 0 & 0 & 0 & 1 & 1 \\ \hline \end{array} \begin{array}{c} \sigma_1 \\ \sigma_2 \\ x_1 \\ y_1 \\ z_1 \\ x_2 \\ y_2 \\ z_2 \\ x \\ y \\ z \end{array} \equiv_2 \begin{array}{c} \begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \\ 0 \end{bmatrix} \\ \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} \end{array}$$

4 Discussion

Work on deriving systems of equalities [9] and inequalities [4] between program variables dates back to the very early days of abstract interpretation. Congruence domains were pioneered by Granger [5,6] who proposed, among other things,

using sets of generators for representing congruence equations and showed that congruence equations satisfied the ascending chain condition.

Recently there has been a resurgence of interest in inferring both linear [7,14] and congruence relationships [1,15], mainly from the perspective of improving efficiency, for instance, by applying randomisation [7], or fusing the domain operations with the fixed-point calculations [14], or refining the conversion between equations and generators [1], or bounding the size of the coefficients in the representation [1,15]. An interesting twist to linear equalities was given by Leroux [13] who has applied the disjunctive closure of this domain in model checking.

Our work revisits congruence analysis, not to enhance efficiency, but to improve precision. Precision is refined by capturing the semantics of a basic block accurately as a system of propositional constraints. These are combined with formulae that express congruence equations that hold upon entry to the block. The constraints that hold at the end of the block are then abstracted as a system of optimal congruence equations. This avoids the need to construct specialised transfer functions for affine assignment, nondeterministic assignment, etc. Instead all primitives, linear and non-linear, can be handled uniformly by translating them into systems of propositional constraints using transformations devised for bounded model checking [3,8,20].

An issue for future work is extending the intra-procedural analysis to inter-procedural analysis and systematic benchmarking. In its present form, the analyser consists of a Prolog and a Java component that are linked with temporary files. The Prolog component translates basic blocks, congruence equalities and congruence disequalities into propositional formulae and then applies CNF conversion [16] to construct a DIMACS file for SAT4J. The Java component implements triangularisation and join. The fixed-point is under manual control since it is both useful and pleasing to watch as the summaries converge onto a loop invariant. However, the Prolog component needs to be extended to translate other operations into formulae in order to deploy the analysis on other code and particularly programs that apply bit-level programming tricks [18].

5 Conclusion

This paper shows how congruence equations, with a modulo that is a power of two, fit elegantly with SAT solving and a relational bit-level encoding of the behaviour of the program, to derive invariants for programs that contain non-linear operations such as bit twiddling. The work calls for further research into methods in which SAT solvers are applied repeatedly to infer abstractions drawn from abstract domains that satisfy the ascending chain condition.

Acknowledgments. This work was funded by EPSRC projects EP/C015517, EP/E033105 and EP/F012896. We thank Paul Docherty for motivating discussions on reverse engineering, Neil Kettle and Axel Simon for their comments on SAT-solving and Gift Nuka for his help with Floyd-style assertions.

References

1. Bagnara, R., Dobson, K., Hill, P.M., Mundell, M., Zaffanella, E.: Grids: A Domain for Analyzing the Distribution of Numerical Values. In: Puebla, G. (ed.) LOPSTR 2006. LNCS, vol. 4407, pp. 219–235. Springer, Heidelberg (2007)
2. Bryant, R.E.: On the complexity of VLSI implementations and graph representations of Boolean functions with application to integer multiplication. *IEEE Transactions on Computers* 40(2), 205–213 (1991)
3. Clarke, E., Kroening, D., Lerda, F.: A tool for checking ANSI-C programs. In: Jensen, K., Podelski, A. (eds.) TACAS 2004. LNCS, vol. 2988, pp. 168–176. Springer, Heidelberg (2004)
4. Cousot, P., Halbwachs, N.: Automatic discovery of linear constraints among variables of a program. In: *Symposium on Principles of Programming Languages*, pp. 84–97. ACM, New York (1978)
5. Granger, P.: Static analysis of arithmetical congruences. *International Journal of Computer Mathematics* 30, 165–190 (1989)
6. Granger, P.: Static analyses of linear congruence equalities among variables of a program. In: Abramsky, S. (ed.) CAAP 1991 and TAPSOFT 1991. LNCS, vol. 493, pp. 167–192. Springer, Heidelberg (1991)
7. Gulwani, S., Necula, G.C.: Discovering affine equalities using random interpretation. In: *Principles of Programming Languages*, pp. 74–84. ACM, New York (2003)
8. Jackson, D., Vaziri, M.: Finding bugs with a constraint solver. In: *International Symposium on Software Testing and Analysis*, pp. 14–25. ACM, New York (2000)
9. Karr, M.: Affine relationships among variables of a program. *Acta Informatica* 6, 133–151 (1976)
10. Kettle, N., King, A., Strzemecki, T.: Widening ROBDDs with prime implicants. In: Hermanns, H., Palsberg, J. (eds.) TACAS 2006 and ETAPS 2006. LNCS, vol. 3920, pp. 105–119. Springer, Heidelberg (2006)
11. King, A., Søndergaard, H.: Inferring congruence equations using SAT. Technical Report 1-08, Computing Laboratory, University of Kent, CT2 7NF (2008)
12. Le Berre, D.: A satisfiability library for Java, <http://www.sat4j.org>
13. Leroux, J.: Disjunctive invariants for numerical systems. In: Wang, F. (ed.) ATVA 2004. LNCS, vol. 3299, pp. 93–107. Springer, Heidelberg (2004)
14. Müller-Olm, M., Seidl, H.: A Note on Karr’s Algorithm. In: Díaz, J., Karhumäki, J., Lepistö, A., Sannella, D. (eds.) ICALP 2004. LNCS, vol. 3142, pp. 1016–1028. Springer, Heidelberg (2004)
15. Müller-Olm, M., Seidl, H.: Analysis of modular arithmetic. *ACM Transactions on Programming Languages and Systems* 29(5) (August 2007) (Article 29)
16. Plaisted, D.A., Greenbaum, S.: A structure-preserving clause form translation. *Journal of Symbolic Computation* 2(3), 293–304 (1986)
17. Ravi, K., McMillan, K.L., Shiple, T.R., Somenzi, F.: Approximation and decomposition of binary decision diagrams. In: *Design Automation Conference*, pp. 445–450. IEEE Press, Los Alamitos (1998)
18. Warren Jr., H.S.: *Hacker’s Delight*. Addison Wesley, Reading (2003)
19. Wegner, P.: A technique for counting ones in a binary computer. *Communications of the ACM* 3(5), 322–322 (1960)
20. Xie, Y., Aiken, A.: SATURN: A scalable framework for error detection using Boolean satisfiability. *ACM Transactions on Programming Languages and Systems* 29(3) (2007) (Article 16)