

InvGen: An Efficient Invariant Generator

Ashutosh Gupta and Andrey Rybalchenko

Max Planck Institute for Software Systems (MPI-SWS)

Abstract. In this paper we present INVGEN, an automatic linear arithmetic invariant generator for imperative programs. INVGEN’s unique feature is in its use of dynamic analysis to make invariant generation order of magnitude more efficient.

1 Introduction

Program verification relies on invariants for reasoning about sets of reachable states [3]. Synthesizing invariants that satisfy a given assertion is a difficult task. The scalability of existing approaches to invariant generation is severely limited due to the high computation cost of the underlying symbolic reasoning techniques.

We present INVGEN, an automatic tool for the generation of linear arithmetic invariants for imperative programs. INVGEN uses a constraint-based approach to generate invariants [2]. INVGEN combines it with static and dynamic analysis techniques to solve constraints efficiently [4]. INVGEN provides an order of magnitude efficiency improvement wrt. the existing tools, see [5]. This improvement enables INVGEN’s application for automatic software verification, e.g., for refinement of predicate abstraction by generating invariants for program fragments determined by spurious counterexamples [1].

In this paper, we describe the design and implementation of INVGEN and present a collection of optimizations that were necessary to achieve the desired scalability. We also describe our experience with applying INVGEN on challenge benchmarks [8] and micro-benchmarks, which containing code fragments that are difficult to analyze automatically. The experimental results indicate that INVGEN is a practical tool for software verification.

2 Usage

INVGEN takes as input a program over linear arithmetic expressions that is given in C syntax or as a transition relation.¹ INVGEN uses a template based technique to compute invariants, which requires that a set of templates is given as input to the tool. A template is a Boolean combination of linear inequalities over program variables with coefficients that are kept as parameters. As a result,

¹ See [5] for the syntax of transition relations.

```

1: x=0;
2: assume(n>0);
3: while(x<n){
4:   x++;
5: }
6: assert(x==n);

```

Fig. 1. Example `simple.c`.

```

1: x=0;
2: while(x<n){
3:   x++;
4: }
5: if(n>0)
6:   assert(x==n);

```

(a)

```

1: assume(p);
2: while(p){
3:   ...
4:   assume(q);
5: }
6: assert(q);

```

(b)

Fig. 2. Examples requiring disjunctive invariants.

INVGEN either returns an invariant that proves the non-reachability of the error location or fails.

INVGEN offers a collection of heuristics to improve scalability of invariant generation, and allows one to explore various combination of these heuristics. Support of dynamic/static analysis can be disabled/enabled with some parameters. Dynamic analysis can be performed using concrete or symbolic execution. Other tools can be used to generate test executions for INVGEN. INVGEN can run in the server mode, in which it communicates its input/output via standard input and output streams. The full set of options is described online [5].

3 Tool

In this section we briefly present the algorithm used by INVGEN and focus on the tool design and implementation.

3.1 Algorithm

INVGEN computes linear arithmetic invariants that prove the non-reachability of the error location. We will call such invariants *safe*. INVGEN applies the template based approach [2] for the invariant generation, which assumes an *invariant template* at each cut-point location in the program, i.e., at loop entry locations. Each invariant template consists of parameterized linear inequalities over the program variables. The specific goal of invariant generation is to find an instantiation of the template parameters that yields a safe invariant. The invariant templates at the start and error locations of the program are `true` and `false`, respectively. A template instantiation yields an invariant if each program path `stmt1;...;stmtN`; between each pair ℓ and ℓ' of adjacent cut-points satisfies the *inductiveness condition*: the instantiated invariant templates $\eta(\ell)$ and $\eta(\ell')$ at the respective cut-points together with the program path form a valid Hoare triple $\{\eta(\ell)\}\text{stmt1};\dots;\text{stmtN};\{\eta(\ell')\}$.² We translate the inductiveness condition into an arithmetic constraint over the template parameters. Each solution of this constraint yields a safe invariant.

² $\{P\}\text{stmts}\{Q\}$ is a valid Hoare triple if each computation that starts in a state satisfying the assertion P either terminates in a state satisfying Q or diverges [7]. In our case, program paths are loop free and hence terminating.

See Figure 1. The example program `simple.c` contains a statement `assert(p)`, which is a short form for `if (!p) ERROR;`. The entry location of the `while` loop is a cut-point location ℓ_3 . At this location we assume an invariant template $\alpha_x x + \alpha_n n \leq \alpha$, where $\alpha_x, \alpha_n, \alpha$ are unknown parameters. The inductiveness condition for the path from the start location to the loop entry requires the validity of the triple $\{\text{true}\}x = 0; \text{assume}(n > 0); \{\eta(\ell_3)\}$, the condition for the loop iteration is $\{\eta(\ell_3)\}\text{assume}(x < n); x = x + 1; \{\eta(\ell_3)\}$, and the path from the loop entry to the error location produces $\{\eta(\ell_3)\}\text{assume}(x \geq n); \text{assume}(x \neq n); \{\text{false}\}$. The resulting arithmetic constraint has a following solution for template parameters: $\alpha_x = 1, \alpha_n = -1$, and $\alpha = 0$. Hence, $x - n \leq 0$ is a safe invariant.

Arithmetic constraints that encode the inductiveness condition contain non-linear arithmetic expressions and are difficult to solve. INVGEN improves performance of the constraint solving by adding information collected using dynamic and static analysis techniques.

In the dynamic analysis phase, we execute the program according to a coverage criterion and collect the reached states. By definition, each reached state must satisfy program invariants. For each reached state, we substitute the program variables occurring in the invariant template at the corresponding control location by their values in the state. Thus, we obtain linear constraints over template parameters. For `simple.c`, assume a reached state at the loop entry such that $(x = 0, n = 1)$. After substituting program variable by their values in the above template for the location ℓ_3 we obtain the constraint over template parameters $\alpha_x 0 + \alpha_n 1 \leq \alpha$. Such linear constraints improve the efficiency of constraint solving [4].

In the static analysis phase, we first apply abstract interpretation in a eager fashion to compute a collection of invariants – which are not necessarily strong enough to prove the non-reachability of the error location – that holds for the program, and then annotate the program with these invariants. The invariants obtain using abstract interpretation allow INVGEN to focus on the synthesis missing, complex invariants in a goal directed way. Such invariants may be missed by the abstract interpreter due to the necessary precision loss of join and widening operators and the limited precision of efficient abstract domains.

3.2 Design

We present the design of INVGEN in Figure 3. The input program is passed to the dynamic and static analyzers. The results of each analysis together with the program and the templates are passed to the constraint generator. The generated constraints are solved by a constraint solver. If the solver succeeds then INVGEN returns a safe invariants.

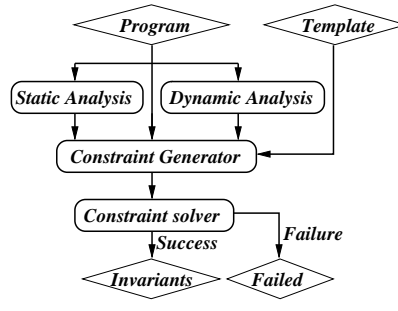


Fig. 3. INVGEN design.

3.3 Implementation

Figure 4 outlines the implementation of INVGEN. It is divided into two executables, *frontend* and INVGEN.

The frontend executable contains a CIL [10] based interface to C and an abstract interpreter INTERPROC [9]. The frontend takes a program procedure written in C language as an input, and applies INTERPROC on the program three times using the interval, octagon, and polyhedral abstract domains. Then, the frontend outputs the transition relation of the program that is annotated with the results computed by INTERPROC. See [5] for the output format.

Next, we describe the components of INVGEN, following Figure 4.

Program minimizer INVGEN minimizes the transition relation of the program to reduce the complexity of constraint solving. INVGEN computes a minimal set of cut-point locations, and replaces each cut-point free path by a single, compound program transition. The unsatisfiable and redundant transitions are eliminated. At this phase, the invariants obtained from INTERPROC can lead to the elimination of additional transitions.

Dynamic analysis INVGEN collects dynamic information for the minimized program using either concrete and symbolic execution. In case of concrete execution, INVGEN collects a finite set of reachable states by using a guided testing technique. Otherwise, INVGEN performs a bounded, exhaustive symbolic execution of the program. By default, the bound is set to the number of cut-points in the program. The user can limit the maximum number of visits for each cut-point during the symbolic execution.

Simplifier INVGEN simplifies all arithmetic constraints locally at each step of the algorithm.

Consider the phase when the dynamic analysis using symbolic execution produces additional linear constraints over template parameters. These constraints contain existentially quantified variables, and the scope of the quantifier is rather limited. For `simple.c`, assume that a symbolic state $\mathbf{x} = 0 \wedge \mathbf{n} \geq 0$ is reached at the loop entry location. For any template evaluation, the symbolic state is subsumed by the invariant, hence $(\mathbf{x} = 0 \wedge \mathbf{n} \geq 0) \rightarrow (\alpha_x \mathbf{x} + \alpha_n \mathbf{n} \leq \alpha)$. After the elimination of program variables INVGEN obtains the constraints $\lambda \geq 0 \wedge -\lambda = \alpha_n \wedge \alpha \geq 0$, where λ is an existentially quantified variable. λ does not appear anywhere in the constraint and can be eliminated locally. As a result, we obtain $\alpha_n \leq 0 \wedge \alpha \geq 0$.

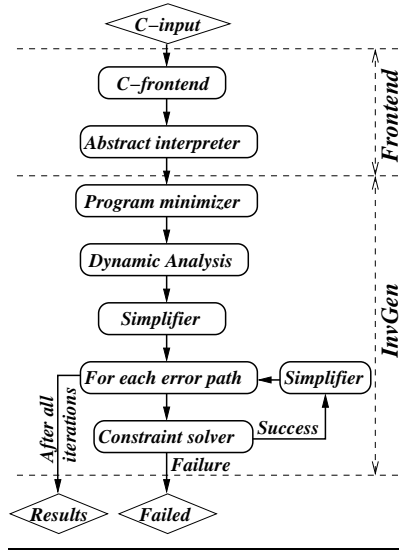


Fig. 4. INVGEN implementation.

INVGEN also simplifies constraints obtained by the concrete execution and abstract interpretation.

Constraint solver The inductiveness conditions result in non-linear arithmetic constraints. For `simple.c`, the inductiveness condition for the error path `{η}assume(x > n); {false}` translates to the implication $\alpha_x x + \alpha_n n \leq \alpha \wedge x > n \rightarrow 1 \leq 0$ with universal quantification over the program variables. After the elimination of the program variables, INVGEN obtains the non-linear constraint $\lambda \geq 0 \wedge \delta \geq 0 \wedge \lambda \alpha_x = \delta \wedge \lambda \alpha_n + \delta = 0 \wedge \lambda \alpha + 1 \leq \delta$, where λ and δ are existentially quantified, non-negative variables that encode the implication validity. In practice, such existentially quantified variables range over a small domain, typically they are either 0 or 1.

INVGEN leverages this observation in order to solve the constraints by performing a case analysis on the variables with small domain. Each instance of case analysis results in a linear constraint over template parameters, which can be solved using a linear constraint solver. This approach is incomplete, since INVGEN does not take all possible values during the case analysis, however it is effective in practice.

Multiple paths to error location A program may have multiple cut-point free paths that lead to the error location, which we refer to as error paths. INVGEN deals with multiple error paths in an incremental fashion for efficiency reasons. Instead taking inductiveness conditions for all error paths into account, INVGEN computes a safe invariant for one error path at a time. Already computed invariants are used as strengthening when dealing with the remaining error paths.

4 Discussion

We highlight several important aspects of INVGEN in detail.

Abstract interpretation The template instantiation computed by INVGEN may rely on the invariants discovered during the abstract interpretation phase. INVGEN keeps track of such invariants and reports them to the user.

Template generation The applicability of INVGEN depends on the choice of invariant templates. If the template is too expressive, e.g., if it admits a number of conjuncts that is larger than required, then the efficiency of INVGEN is decreasing due to the increased difficulty of constraint solving. A template that is not expressive enough cannot yield a desired invariant.

In our experience, the majority of invariants require a template that is a conjunction of two linear inequalities. This surprisingly small number of conjuncts is due to the strengthening using invariants obtained during the abstraction interpretation phase—INVGEN focuses on the missing invariants. By default, for each cut-point location INVGEN assumes a conjunction of two parametric linear inequalities as a template.

File	INVGEN**	INVGEN*	INVGEN
Seq	23.0s	0.5s	0.5s
Seq-z3	23.0s	0.5s	0.5s
Seq-len	T/O	T/O	2.8s
nested	T/O	17.0s	2.3s
svd(light)	T/O	10.6s	14.2s
heapsort	T/O	19.2s	13.3s
mergesort	T/O	142s	170s
Spam-loop	T/O	0.28s	0.4s
apache-tag	0.4s	0.6s	0.7s
sendmail-qp	0.3s	0.3s	0.3s

Table 1. Performance of INVGEN on benchmark inspired by [8]. INVGEN* does not apply dynamic analysis. INVGEN** applies neither static nor dynamic analysis. “T/O” means time out after 10 minutes.

File	BLAST	BLAST + INVGEN
Seq	diverge	8
Seq-len	diverge	9
fregtest	diverge	3
sendmail-qp	diverge	10
svd(light)	144	43
Spam-loop	51	24
apache-escape	26	20
apache-tag	23	15
sendmail-angle	19	15
sendmail-7to8	16	13

Table 2. Application of INVGEN for the predicate discovery in BLAST using path invariants. We show the number of refinement steps required to prove the property.

Disjunction The user can provide disjunctive templates to INVGEN, which yields disjunctive invariants. Currently, INVGEN is not practical for disjunctive invariants.

Nevertheless, programs that require disjunctive invariants appear often in practice, as illustrated in Figure 2. Example 2(a) relies on a disjunctive loop invariant $n \geq x \vee n < 0$ to prove the assertion. Similarly, Example 2(b) requires a disjunctive loop invariant $p \vee q$.

Avoiding such program patterns improves the effectiveness of INVGEN. For example, adding a commonly used assumption $n \geq 0$ at the start location of the program in Figure 2(a) eliminates the need for disjunctive invariants.

Concrete vs. symbolic dynamic analysis In our evaluation, dynamic analysis using symbolic execution performs better than with the concrete execution. Symbolic execution produces a smaller set of constraints and leads to the improvement of the constraint solving efficiency that is never worse the one achieved by the concrete execution.

5 Experiences

Verification benchmarks We applied INVGEN on a collection of programs that are difficult for state-of-the-art software verification tools [8]. The collection consists of 12 programs. Due to short running times, we present the aggregated data and do not provide information for each program individually. Using the polyhedral abstract domain, INTERPROC computes invariants that are strong enough to prove the assertion for six programs in the collection. INVGEN handles 11 examples in 6.3 seconds, and times out on one program.

Effect of static and dynamic analysis The collection [8] does not allow us to perform a thorough experimental evaluation of INVGEN, since the running

times on these examples are too short. We constructed a set of more interesting programs that is inspired by [8] by extending the programs from this collection with additional loops and branching statements. Table 1 shows the effect of static and dynamic analysis when dealing with the constructed examples.

Random input We developed a tool for the generation of random program. We try these random programs with INVGEN without testing then with testing feature. We generated 17 programs with three cut-point and one error location in each of them, see [5]. We observed that for these programs, dynamic analysis either improves or at least does not decrease the efficiency of invariant generation. In four cases, the efficiency of INVGEN increased by two orders of magnitude.

Integration with Blast We have modified the abstraction refinement procedure of the BLAST software model checker [6] by adding predicate discovery using path invariants [1]. Table 2 indicates that constraint based invariant generation can be an effective tool for refining predicate abstraction. For several examples, BLAST diverged due to the disability to find the right set of predicates. In these cases, INVGEN enabled the successful termination of the verification attempt, while for other cases it reduced the number of the number of refinement iterations by 25–400%.

References

1. D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Path invariants. In *Proc. PLDI*, pages 300–309. ACM Press, 2007.
2. M. Colón, S. Sankaranarayanan, and H. Sipma. Linear invariant generation using non-linear constraint solving. In *CAV*, pages 420–432. Springer, 2003.
3. R. W. Floyd. Assigning meanings to programs. In *Mathematical Aspects of Computer Science*. AMS, 1967.
4. A. Gupta, R. Majumdar, and A. Rybalchenko. From tests to proofs. In *Proc. TACAS*. Springer, 2009.
5. A. Gupta and A. Rybalchenko. InvGen: an efficient invariant generator. <http://www.mpi-sws.org/~agupta/invgen/>.
6. T. Henzinger, R. Jhala, R. Majumdar, and K. McMillan. Abstractions from proofs. In *POPL 04: Principles of Programming Languages*, pages 232–244. ACM, 2004.
7. C. A. R. Hoare. An axiomatic basis for computer programming. *Communications of ACM*, 1969.
8. K. Ku, T. Hart, M. Chechik, and D. Lie. A buffer overflow benchmark for software model checkers. In *Proc. ASE*, 2007.
9. G. Lalire, M. Argoud, and B. Jeannet. The interproc analyze. <http://pop-art.inrialpes.fr/people/bjeannet/bjeannet-forge/interproc/index.html>.
10. G. C. Necula, S. McPeak, S. P. Rahul, and W. Weimer. CIL: Intermediate language and tools for analysis and transformation of C programs. In *Proc. CC*, pages 213–228. Springer, 2002.