

# Incremental Verification of Compiler Optimizations<sup>\*</sup>

Grigory Fedukovich<sup>1</sup>, Arie Gurfinkel<sup>2</sup>, and Natasha Sharygina<sup>1</sup>

<sup>1</sup> University of Lugano, Switzerland,  
{grigory.fedyukovich,natasha.sharygina}@usi.ch  
<sup>2</sup> SEI/CMU, USA, arie@cmu.com

**Abstract** Optimizations are widely used along the lifecycle of software. However, proving the equivalence between original and optimized versions is difficult. In this paper, we propose a technique to incrementally verify different versions of a program with respect to a fixed property. We exploit a safety proof of a program given by a safe inductive invariant. For each optimization, such invariants are adapted to be a valid safety proof of the optimized program (if possible). The cost of the adaptation depends on the impact of the optimization and is often less than an entire re-verification of the optimized program. We have developed a preliminary implementation of our technique in the context of Software Model Checking. Our evaluation of the technique on different classes of industrial programs and standard LLVM optimizations confirms that the optimized programs can be re-verified efficiently.

## 1 Introduction

Program verification is necessary for building reliable software intensive systems. One challenge in using verification is deciding on the right level of abstraction. On one hand, verifying the source code (or a high-level compiler representation) before optimizations gives meaningful verification results to the user. On the other, verifying the binary (or a low-level compiler representation) after optimizations takes compiler out of the trusted computing base. Our experience with the UFO [1] indicates that verification results at both levels are desired.

There are two common techniques for adapting verification results from an original program  $P$  to an optimized program  $Q$ : (1) complete re-verification of  $Q$ ; (2) establish property preserving equivalence (typically a form of a simulation) between  $P$  and  $Q$ . Re-verification is computationally expensive. Establishing a simulation between  $P$  and  $Q$  often requires manual instrumentation of the compiler which is hard to do and maintain [6]. In this paper, we propose an alternative solution that combines the advantages of the two approaches.

We assume that the original program  $P$  comes with a property  $G$ , and that  $P$  satisfies  $G$  (i.e.,  $P \models G$ ). Instead of showing an equivalence between  $P$  and  $Q$ ,

---

<sup>\*</sup> This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense. This material has been approved for public release and unlimited distribution. DM-0000784

we show that  $Q$  satisfies  $G$ . First, by adapting the proof of  $P \models G$ , given by an inductive invariant, to  $Q$ , and then strengthening it by re-verification as needed. Our technique can be seen as a *property-specific equivalence*:  $P$  and  $Q$  are equivalent iff they both satisfy  $G$ .

We evaluate our approach on the `instcombine` optimization of LLVM that does local optimizations (such as turning  $x = 1 + 1$  into  $x = 2$ ). Our experiments show that the approach is very effective. In many cases, the complete safety proof can be transferred between the original and the optimized programs. Whenever re-verification was required, it was insignificant.

## 2 From Optimization to Evolution

In this section, we formally define the problem of *incremental property-directed verification of optimizations*. We begin with formal definitions of programs, safety proofs, and admissible optimizations. A “large-block” representation of a *program* is a tuple  $P = (V, en, err, E, \tau_P)$ , where  $V$  is a set of cutpoints (i.e., locations which represent heads of some loops);  $en, err \in V$  are designated entry and error locations, respectively;  $E \subseteq V \times V$  is the control-flow relation (represent a loop-free program fragments),  $\tau_P : E \rightarrow Stmt^*$  maps control edges to loop-free program fragments. An example of a program is shown in Fig. 1a. We call the graph  $(V, E)$  the Cut-Point Graph (CPG) which collapses the more fine grained Control-Flow Graph (CFG).

We write  $\vdash X$  to mean that  $X$  is valid. Let  $Expr$  be a set of expressions over program variables,  $pre, post \in Expr$  and  $S$  a loop-free program fragment. We write  $\vdash \{pre\} S \{post\}$  to mean that  $pre$  and  $post$  are pre- and post-conditions of  $S$ , i.e., whenever  $S$  starts in a state satisfying  $pre$ , if  $S$  terminates, it ends in a state satisfying  $post$ . A *safety proof* of  $P$  is a mapping  $\psi : V \rightarrow Expr$  such that

$$\forall (u, v) \in E \cdot \vdash \{\psi(u)\} \tau_P(u, v) \{\psi(v)\} \quad \psi(err) \rightarrow \perp \quad (1)$$

An *optimization* of  $P$  is a program  $Q = (V, en, err, E, \tau_Q)$ , that differs from  $P$  only in the labeling of the edges of the CPG. Note that this definition limits optimizations to changes that do not affect the loop-structure of  $P$  (i.e., loops cannot be added or removed). However, since we deal with “large-blocks”, loop

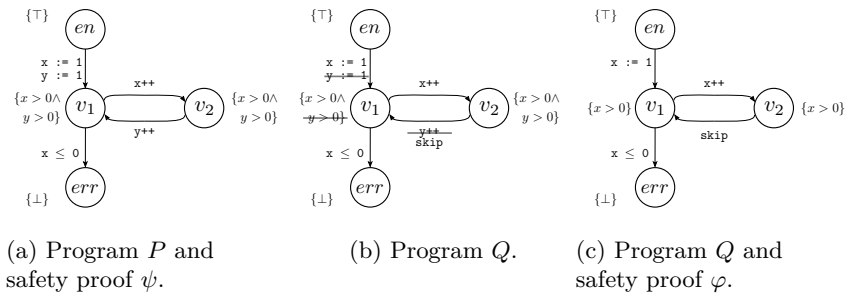


Figure 1: A simple program and an optimization.

---

**Algorithm 1:** OPTVERIFY( $P, \psi, Q$ )

**Input:** *Original* program  $P$ , safety proof  $\psi$  of  $P$ , *new* program  $Q$

**Output:** A pair  $\langle res, \varphi \rangle$  s.t.  $res \rightarrow \varphi$  is safety proof of  $Q$

- 1  $\sigma \leftarrow \text{GUESSMAP}(P, Q)$  ▷ Guess a map between variables of  $P$  and  $Q$
  - 2  $\pi \leftarrow \text{MKIND}(\psi\sigma, Q, P)$  ▷ Weaken candidate  $\psi\sigma$  until it is inductive for  $Q$
  - 3  $res, \varphi \leftarrow \text{VERIFY}(Q, \pi)$  ▷ Strengthen  $\varphi$  until it is SAFE for  $Q$
  - 4 **return**  $\langle res, \varphi \rangle$
- 

unrolling is admissible. At the same time, we do not put any other restrictions on  $Q$ . In particular, we do not require for  $Q$  to simulate  $P$ . An example of optimization  $Q$  of  $P$  is shown at the Fig. 1b in which the variable  $y$  was removed

The problem of incremental property-directed verification is: Given a program  $P$ , a safety proof  $\psi$  of  $P$ , and an optimization  $Q$  of  $P$ , *adapt* a safety proof  $\psi$  to a safety proof  $\varphi$  of  $Q$ , or show that  $Q$  does not admit a safety proof (i.e., has a counterexample).

Thus, we view the problem as a variant of *upgrade checking* (as opposed to verification). The key concept in upgrade checking is adapting verification results from one program to another. We extend this concept to safety proofs. Let  $P$ ,  $\psi$ , and  $Q$  be as above. An adaptation  $\pi$  of  $\varphi$  to  $Q$  is a weakening of  $\psi$  such that

$$\forall w \in V \cdot \psi(w) \rightarrow \pi(w) \tag{2}$$

$$\forall (u, v) \in E \cdot \vdash \{ \pi(u) \} \tau_Q(u, v) \{ \pi(v) \} \tag{3}$$

Intuitively, (2) says that  $\pi$  is a weakening of  $\psi$ , and (3) says that  $\pi$  is inductive. Note that if  $\pi(err) = \perp$ , then  $\pi$  is a safety proof of  $Q$ . Otherwise,  $\pi$  is simply an inductive invariant.

### 3 Our Solution

OPTVERIFY is shown in Alg. 1. The input is an *original* program  $P = (V, en, err, E, \tau_P)$ , a safety proof  $\psi : V \rightarrow Expr$  of  $P$  and a *new* program  $Q = (V, en, err, E, \tau_Q)$ . The output is a safety proof  $\varphi$  of  $Q$  or a counter-example. We require that  $P$  and  $Q$  share the CPG (i.e., same loop-heads), and differ only in the labeling of edges.

OPTVERIFY consists of three steps: (1) a mapping  $\sigma$  between the variables of  $P$  and  $Q$  (line 1) is guessed and is used to syntactically transfer the safety proof  $\psi$  to a candidate proof  $\psi\sigma$  of  $Q$ . (2)  $\psi\sigma$  is weakened to  $\pi$  using MKIND until  $\pi$  is an inductive invariant (line 2). At this point,  $\pi$  is inductive, but not safe. (3)  $\pi$  is strengthened if necessary (and possible) to  $\varphi$  by a standalone verifier (line 3).

In our implementation,  $\sigma$  is guessed by syntactically matching variable names: a variable  $v$  of  $P$  is mapped to a variable  $v$  of  $Q$  if  $v$  exists in  $Q$ , and to a fresh symbol otherwise. To implement VERIFY, we use UFO [1] that iteratively strengthens the given inductive invariant. In particular, if  $\pi$  is already safe, UFO returns immediately and  $\varphi \equiv \pi$ . A verifier that cannot work by strengthening a given invariant would be useless as it would amount to verifying  $Q$  from scratch.

---

**Algorithm 2:** MKIND( $\psi, Q, P$ )**Input:** Candidate invariant  $\psi$ , *new* program  $Q$ , *original* program  $P$ **Output:** Inductive invariant  $\pi : V \rightarrow Expr$  of  $Q$ 

```
1  $\pi \leftarrow \psi$ ;  $W \leftarrow \{(u, v) \in E \mid \tau_P(e) \neq \tau_Q(e)\}$ ;  
2 while  $W \neq \emptyset$  do  
3    $(u, v) \leftarrow \text{GETWTO\textsmaller{SMALLESTEDGE}}(W)$ ;           // according to the WTO  
4    $pre \leftarrow \pi(u)$ ;  $post \leftarrow \pi(v)$ ;  
5   if  $(\vdash \{pre\} \tau_Q(u, v) \{post\})$  then  $W \leftarrow W \setminus \{(u, v)\}$ ; // use SMT-solver  
6   else  
7      $\pi(v) \leftarrow \text{WEAKPOST}(pre, \tau_Q(u, v), post)$ ;           // see Alg 3  
8      $W \leftarrow (W \setminus \{(u, v)\}) \cup \{(v, x) \in E \mid x \in V\}$   
9 return  $\pi$ 
```

---

Our key contribution is an algorithm MKIND for weakening a candidate invariant. We describe it in details in the rest of this section.

MKIND is shown in Alg. 2. The input is a candidate invariant  $\psi$ , a *new* program  $Q$  and an *original* program  $P$ . The output is an inductive weakening  $\pi$  of  $\psi$  (i.e.,  $\pi$  satisfies (2) and (3)) for  $Q$ . MKIND maintains a work-list  $W \subseteq E$  that is initialized with all the edges  $(u, v) \in E$  on which  $P$  and  $Q$  disagree (i.e.,  $\tau_P(u, v) \neq \tau_Q(u, v)$ ). In each iteration of the main loop, first, an edge  $(u, v) \in W$  that is least in the Weak Topological Ordering (WTO) [2] (in which inner loops are traversed before outer loops) is picked. Second, an SMT-solver is used to check whether the current values of  $\pi(u)$  and  $\pi(v)$  form a valid Hoare-triple for the corresponding loop-free code fragment  $\tau_Q(u, v)$ . If this is not the case, the post-condition  $\pi(v)$  is weakened until the triple becomes valid and all outgoing edges of  $v$  are added to  $W$ . Soundness of MKIND is immediate – the work-list is empty only if every edge is annotated with a valid pre- and post-condition pair. Termination follows from the fact that at each iteration either the work-list is reduced, or a post-condition is weakened, and, our implementation of WEAKPOST allows only for finitely many weakening steps.

WEAKPOST is shown in Alg. 3. The input is a pre-condition  $pre$ , a post-condition  $post$  and a loop-free program fragment  $S$ . The output is a weakening  $post'$  of  $post$  such that  $\{pre\} S \{post'\}$  is valid. We assume that the post-condition  $post$  is given as a conjunction of lemmas, i.e.,  $post = \bigwedge_i \ell_i$ . The algorithm computes the (possibly empty) subset of  $\{\ell_i\}$  that forms a valid post-condition.

The naive implementation of WEAKPOST iteratively checks whether each  $\ell_i$  is a post-condition. Instead, we use an incremental SMT solver to do this enumeration efficiently. We assume that in addition to the SMTSOLVE API, an SMT solver has the method SMTASSERT to add constraints to the current context. First, we compute an SMT-formula that encodes the verification condition (VC) of  $S$ . We use helper method MKVC that implements VC-generation from [5]. Second, we construct  $vc$  that determines the validity of pre- and post-condition by conjoining the pre-condition and negation of the post-condition (line 2) to the VC. Note that the post-condition is asserted under assumptions, encoded by Boolean variables  $x_i$  such that lemma  $\ell_i$  is active iff  $x_i$  is true. We then itera-

---

**Algorithm 3:** WEAKPOST( $pre, S, post$ )

**Input:**  $pre, post \in Expr$ ;  $post = \bigwedge_{i=0}^n \ell_i$ ;  $S \in Stmt^*$   
**Output:**  $post' \in Expr$ , such that  $\vdash \{pre\} S \{post'\}$  is valid

- 1 let  $\{x_i \mid 0 \leq i \leq n\}$  be fresh Boolean variables;  $U \leftarrow \{0, \dots, n\}$ ;
- 2  $vc \leftarrow pre \wedge \text{MKVC}(S) \wedge \neg(\bigwedge_{i=0}^n (x_i \rightarrow \ell_i))$ ;
- 3 SMTASSERT( $vc$ );
- 4 **while** SMTSOLVE() = SAT **do**
- 5      $M \leftarrow \text{SMTMODEL}()$ ;
- 6     **foreach**  $\{0 \leq i \leq n \mid M \models x_i\}$  **do** SMTASSERT( $\neg x_i$ );  $U \leftarrow U \setminus \{i\}$ ;
- 7  $post' \leftarrow \bigwedge \{\ell_i \mid i \in U\}$ ;
- 8 **return**  $post'$

---

tively check the validity of  $vc$ . In each iteration, if  $vc$  is satisfiable, we assert  $\neg x_i$  to disable the corresponding lemma(s). This terminates eventually since there are finitely many lemmas and at least one is disabled at every iteration. The conjunction of all active lemmas is returned as  $post'$ .

To summarize we illustrate MKIND on an example from Figs. 1a-1b. Given  $P$ , its safety proof  $\psi$  and  $Q$ , MKIND first checks the validity of the edge  $en \rightarrow v_1$ , i.e.,  $\{\top\} \mathbf{x} := 1 \{x > 0 \wedge y > 0\}$ . It is clearly invalid, thus WEAKPOST is used to weaken  $\pi(v_1)$  to  $x > 0$ . Next, the validity of the triple for the edge  $v_1 \rightarrow v_2$  is checked, i.e.,  $\{x > 0\} \mathbf{x}++ \{x > 0 \wedge y > 0\}$ . Note that the post-condition constructed in the previous step is now used as the pre-condition. Again, the triple is invalid and the post-condition is weakened to  $x > 0$ . Finally, the edges  $v_2 \rightarrow v_1$  and  $v_1 \rightarrow err$  are checked, and MKIND terminates. The final inductive invariant  $\pi$  of  $Q$  is shown in Fig. 1c. In this case,  $\pi$  is also a safety proof ( $\pi(err) \rightarrow \perp$ ), and, therefore,  $\varphi \equiv \pi$ ), so no further strengthening is required.

## 4 Evaluation

We have implemented OPTVERIFY in the UFO framework, and evaluated it on the Software Verification Competition (SVCOMP'13) benchmarks and `instcombine` optimization of LLVM. For each benchmark (300 - 5000 lines of source code), we measured the verification time `oV`, time of OPTVERIFY (MKIND + VERIFY), and time to re-verify from scratch (`uV`). Out of 397 safe benchmarks, we chose 108 that had non-trivial original verification time (`oV` > 1s). Due to lack of space, results are available at <http://www.inf.usi.ch/phd/fedyukovich/niagara>.

In all but 6 cases, re-verification (VERIFY) was insignificant (< 1s). Furthermore, in 67 cases the candidate invariant was already safe and no strengthening was needed. In 52 cases, OPTVERIFY was at least an order of magnitude faster than re-verification (`uV`). This highlights the benefit of the approach. In 34 cases, `instcombine` dramatically reduced verification time (from minutes to < 1s). This shows that sometimes it is better to verify the optimized program first and then adapt the results to the original one. While this can be done using OPTVERIFY, we have not done so yet.

## 5 Related Work

There is a large body of related work on incremental verification, verifying compilers, and translation validation. Here, we only survey some of the most relevant techniques. Our algorithm is inspired by recent advancements in upgrade checking [3] and witnessing compiler transformations [6].

The goal of upgrade checking is to adapt a verification result from an original program  $P$  to an *upgraded* program  $Q$ , where  $Q$  is obtained from  $P$  by changing some of its functions. Current techniques [3] apply to bounded programs with functions, and work by adapting function summaries (i.e., relationship between function’s inputs and outputs) from  $P$  to  $Q$ . In contrast, we work on unbounded programs without functions and adapt safety proofs (i.e., safe inductive invariants). Extending our approach to deal with functions is an interesting direction for future work.

The goal of witnessing compiler transformations is to formally establish equivalence (namely, stuttering simulation) between an original program  $P$  and its optimization  $Q$ . Current technique [6] works by instrumenting every optimization pass to output a witness relating its input and output. In contrast, we are interested in an automated solution. In that, our work is more similar to translation validation [7]. However, we only require equivalence with respect to a given property. At the technical level, our use of CPGs makes the approach less sensitive to many restructuring optimizations such as loop unrolling. We expect that in practice the two approaches can be combined, extending our algorithm, for example, to handle optimizations changing CPG structure.

At its core, our solution is similar to Houdini [4] inference algorithm that constructs inductive invariant out of a set of candidate formulas. However, there are many technical differences. One being our use of incremental SMT solver to speed up the search. More importantly, we need to adapt candidates from one program to another. Since based on discarding some candidate lemmas, our current adaptation strategy is rather rough and might be replaced by a more accurate one which weakens formulas on logical level and guided by results on analysis of concrete optimizations. This remains a challenge for future work.

## References

1. Albarghouthi, A., Gurfinkel, A., Chechik, M.: UFO: A Framework for Abstraction- and Interpolation-Based Software Verification. In: CAV. LNCS, vol. 7358, pp. 672–678. Springer (2012)
2. Bourdoncle, F.A.: Efficient Chaotic Iteration Strategies with Widening. In: FMPA. LNCS, vol. 735, pp. 128–141. Springer (1993)
3. Fedyukovich, G., Sery, O., Sharygina, N.: eVolCheck: Incremental Upgrade Checker for C. In: TACAS. LNCS, vol. 7795, pp. 292–307. Springer (2013)
4. Flanagan, C., Leino, K.R.M.: Houdini: an Annotation Assistant for ESC/Java. In: FME. LNCS, vol. 2021, pp. 500–517. Springer (2001)
5. Gurfinkel, A., Chaki, S., Sapra, S.: Efficient Predicate Abstraction of Program Summaries. In: NFM. LNCS, vol. 6617, pp. 131–145. Springer (2011)

6. Namjoshi, K.S., Zuck, L.D.: Witnessing program transformations. In: SAS. LNCS, vol. 7935, pp. 304–323. Springer (2013)
7. Necula, G.C.: Translation validation for an optimizing compiler. In: PLDI. pp. 83–94. ACM (2000)

## A Experimental results

benchmark (name)	oV (sec)	optVerify			speedup (X)	uV (sec)
		mkInd (sec)	proof candidate	Verify (sec)		
s3_srvr.blast.15_safe.i.cil.o3.bc	248.2	0.23	safe	0.06	∞	timeout
s3_srvr.7_safe.cil.o3.bc	324	0.25	safe	0.07	∞	timeout
s3_srvr.2_safe.cil.o3.bc	281.95	0.24	safe	0.08	∞	timeout
s3_srvr.3_safe.cil.o3.bc	398.65	0.27	safe	0.1	∞	timeout
s3_srvr.blast.13_safe.i.cil.o3.bc	270.63	0.44	safe	0.12	∞	timeout
token_ring.08_safe.cil.o3.bc	322.45	0.6	safe	0.18	∞	timeout
pc_stlfo.2_safe.cil.o3.bc	308.8	0.51	safe	0.22	∞	timeout
token_ring.09_safe.cil.o3.bc	468.19	0.82	safe	0.28	∞	timeout
s3_srvr.8_safe.cil.o3.bc	243.41	0.2	safe	0.06	1450.23	377.06
s3_srvr.blast.16_safe.i.cil.o3.bc	60.04	0.12	safe	0.04	953	152.48
s3_srvr.blast.10_safe.i.cil.o3.bc	77.82	0.18	safe	0.04	833.5	183.37
s3_srvr.blast.11_safe.i.cil.o3.bc	280.88	0.28	safe	0.08	715.11	257.44
s3_srvr.blast.16_safe.i.cil.o3.bc	14.92	0.1	safe	0.04	571.57	80.02
s3_srvr.blast.06_safe.i.cil.o3.bc	181.96	0.2	safe	0.06	499.96	129.99
s3_clnt.3_safe.cil.o3.bc	94.88	0.23	safe	0.08	401	124.31
s3_srvr.blast.01_safe.i.cil.o3.bc	320.2	0.25	safe	0.06	375.84	116.51
s3_srvr.blast.09_safe.i.cil.o3.bc	92.75	0.25	safe	0.06	361.23	111.98
token_ring.06_safe.cil.o3.bc	136.59	0.37	safe	0.12	346.78	169.92
s3_srvr.blast.15_safe.i.cil.o3.bc	22.69	0.13	safe	0.04	337.88	57.44
token_ring.07_safe.cil.o3.bc	159.64	0.48	safe	0.16	301.69	193.08
s3_srvr.blast.02_safe.i.cil.o3.bc	54.64	0.16	safe	0.04	294.3	58.86
s3_srvr.blast.12_safe.i.cil.o3.bc	19.75	0.14	safe	0.05	257.37	48.9
s3_srvr.blast.13_safe.i.cil.o3.bc	91.51	0.26	safe	0.06	256.13	81.96
s3_srvr.blast.11_safe.i.cil.o3.bc	52.65	0.22	safe	0.06	256.04	71.69
s3_clnt.blast.03_safe.i.cil.o3.bc	75.36	0.28	safe	0.08	249.61	89.86
token_ring.06_safe.cil.o3.bc	265.38	1.18	safe	0.18	231.2	314.43
s3_srvr.blast.07_safe.i.cil.o3.bc	57.76	0.22	safe	0.08	221.7	66.51
token_ring.07_safe.cil.o3.bc	262.47	1.22	safe	0.14	221.35	301.04
s3_clnt.2_safe.cil.o3.bc	49.67	0.2	safe	0.06	219.38	57.04
s3_clnt.1_safe.cil.o3.bc	32.6	0.15	safe	0.04	203.47	38.66
s3_srvr.4_safe.cil.o3.bc	19.05	0.08	safe	0.03	197.55	21.73
s3_srvr.blast.14_safe.i.cil.o3.bc	44.75	0.15	safe	0.04	191.47	36.38
s3_srvr.3_safe.cil.o3.bc	19.91	0.11	safe	0.04	182.93	27.44
s3_srvr.blast.12_safe.i.cil.o3.bc	23.13	0.12	safe	0.03	173.4	26.01
s3_clnt.blast.02_safe.i.cil.o3.bc	41.78	0.23	safe	0.06	167.52	48.58
s3_srvr.blast.01_safe.i.cil.o3.bc	13.64	0.11	safe	0.05	160.5	25.68
s3_srvr.blast.14_safe.i.cil.o3.bc	192.93	0.28	safe	0.05	130.91	43.2
s3_clnt.1_safe.cil.o3.bc	11.55	0.13	safe	0.04	129.47	22.01

Figure 2: Verification statistics of oV, OPTVERIFY and uV (part 1 of 3)

benchmark (name)	oV (sec)	optVerify			speedup (X)	uV (sec)
		mkInd (sec)	proof candidate	Verify (sec)		
s3_cint_1_safe.cil.o0.bc	11.55	0.13	safe	0.04	129.47	22.01
s3_srvr.blast.06_safe.i.cil.o0.bc	61.02	0.17	safe	0.05	126.86	27.91
token_ring.05_safe.cil.o3.bc	32.18	0.24	safe	0.07	118.97	36.88
s3_cint.blast.01_safe.i.cil.o3.bc	15.44	0.13	safe	0.04	108.82	18.5
s3_srvr.blast.10_safe.i.cil.o0.bc	16.32	0.12	safe	0.04	88.94	14.23
token_ring.04_safe.cil.o3.bc	16.46	0.18	safe	0.05	81.61	18.77
s3_cint_4_safe.cil.o3.bc	12.08	0.13	safe	0.04	81.24	13.81
s3_cint.blast.04_safe.i.cil.o3.bc	12.54	0.14	safe	0.04	80.67	14.52
s3_cint_3_safe.cil.o0.bc	14.13	0.16	safe	0.04	79.05	15.81
s3_srvr.blast.08_safe.i.cil.o0.bc	20.53	0.14	safe	0.04	75.44	13.58
s3_cint.blast.02_safe.i.cil.o0.bc	7.83	0.13	safe	0.03	65	10.4
s3_cint.blast.03_safe.i.cil.o0.bc	9.65	0.14	safe	0.03	64.18	10.91
s3_cint_2_safe.cil.o0.bc	9.03	0.14	safe	0.04	62.94	11.33
s3_cint_4_safe.cil.o0.bc	9.49	0.14	safe	0.04	62.72	11.29
token_ring.05_safe.cil.o0.bc	41.37	0.65	safe	0.1	60.4	45.3
s3_cint.blast.01_safe.i.cil.o0.bc	9.67	0.11	safe	0.03	56.71	7.94
s3_cint.blast.04_safe.i.cil.o0.bc	8.69	0.13	safe	0.03	51.75	8.28
token_ring.04_safe.cil.o0.bc	22.82	0.48	safe	0.08	45.05	25.23
token_ring.03_safe.cil.o3.bc	4.47	0.12	safe	0.04	32.13	5.14
token_ring.03_safe.cil.o0.bc	6.8	0.25	safe	0.05	24.93	7.48
token_ring.02_safe.cil.o3.bc	1.07	0.07	safe	0.02	13.56	1.22
s3_srvr_2_safe.cil.o0.bc	114.89	0.99	safe	0.05	11.42	11.88
pc_sfifo_1_safe.cil.o3.bc	1.36	0.11	safe	0.04	10.47	1.57
token_ring.02_safe.cil.o0.bc	1.34	0.12	safe	0.03	10.2	1.53
pc_sfifo_3_safe.cil.o3.bc	1.32	0.16	safe	0.04	7.4	1.48
32_1_cilled_safe_ok_nondet_linux-3.4-32_1-drivers-- media-dvb-dvb-usb-dvb-usb-vp7045.ko- ldv_main0_sequence_infinite_withcheck_stateful.cil. out.o3.bc	1.72	0.38	safe	0.16	3.96	2.14
32_1_cilled_safe_ok_nondet_linux-3.4-32_1-drivers-- media-dvb-dvb-usb-dvb-usb-vp7045.ko- ldv_main0_sequence_infinite_withcheck_stateful.cil. out.o0.bc	2.72	1.1	safe	0.16	2.48	3.13
elevator_spect14_product03_safe.cil.o3.bc	1.36	0.33	safe	0.08	1.93	0.79
32_1_cilled_safe_ok_nondet_linux-3.4-32_1-drivers-- media-dvb-dvb-usb-dvb-usb-vp7045.ko- ldv_main1_sequence_infinite_withcheck_stateful.cil. out.o3.bc	3.58	1.54	safe	0.54	1.9	3.96
32_1_cilled_safe_ok_nondet_linux-3.4-32_1-drivers-- media-dvb-dvb-usb-dvb-usb-vp7045.ko- ldv_main1_sequence_infinite_withcheck_stateful.cil. out.o0.bc	1.24	2.45	safe	0.45	0.49	1.42
usb_urb-drivers-uwbi1480-dfu-i1480-dfu-usb. ko_safe.cil.out.i.pp.o0.bc	1.63	3	safe	1.78	0.42	2.02
pipeline_safe.cil.o0.bc	88.64	6.7	ind	0.34	16.89	118.94
pipeline_safe.cil.o3.bc	369.05	2.68	ind	2.01	5.56	26.09
s3_srvr_8_safe.cil.o0.bc	67.08	0.31	ind	18.63	2.72	51.57
s3_srvr_7_safe.cil.o0.bc	118.48	0.54	ind	123.16	1.59	196.13
s3_srvr_6_safe.cil.o0.bc	58.44	0.15	ind	48.53	0.47	22.74
s3_srvr_4_safe.cil.o0.bc	37.03	0.19	ind	56.26	0.46	25.7

Figure 3: Verification statistics of oV, OPTVERIFY and uV (part 2 of 3)



benchmark (name)	oV (sec)	optVerify			speedup (X)	uV (sec)
		mkInd (sec)	proof candidate	Verify (sec)		
elevator_spec2_product01_safe.cil.o3.bc	9.54	2.54	ind	0.1	0.09	0.24
elevator_spec2_product17_safe.cil.o3.bc	8.69	2.76	ind	0.08	0.08	0.22
elevator_spec2_product03_safe.cil.o3.bc	8.64	2.62	ind	0.08	0.08	0.22
elevator_spec9_product27_safe.cil.o3.bc	10.34	3.75	ind	0.1	0.06	0.22
elevator_spec2_product19_safe.cil.o3.bc	8.84	4.48	ind	0.08	0.05	0.24
elevator_spec2_product27_safe.cil.o3.bc	10.36	4.03	ind	0.13	0.05	0.19
elevator_spec2_product25_safe.cil.o3.bc	10.7	4.04	ind	0.13	0.05	0.2
elevator_spec9_product09_safe.cil.o3.bc	11.48	4.3	ind	0.14	0.05	0.22
elevator_spec2_product09_safe.cil.o3.bc	10.4	4.13	ind	0.14	0.05	0.23
elevator_spec2_product29_safe.cil.o3.bc	19.62	7.7	ind	0.18	0.05	0.36
elevator_spec9_product25_safe.cil.o3.bc	10.4	3.76	ind	0.18	0.05	0.21
elevator_spec2_product11_safe.cil.o3.bc	8.94	5.18	ind	0.12	0.04	0.22
elevator_spec9_product11_safe.cil.o3.bc	7.71	4.33	ind	0.12	0.04	0.19
elevator_spec2_product31_safe.cil.o3.bc	17.73	6.72	ind	0.16	0.04	0.26
elevator_spec9_product31_safe.cil.o3.bc	26.38	9.72	ind	0.12	0.03	0.28
elevator_spec2_product23_safe.cil.o3.bc	26.54	9.29	ind	0.15	0.03	0.29
elevator_spec2_product21_safe.cil.o3.bc	18.29	9.37	ind	0.15	0.03	0.3
elevator_spec1_product25_safe.cil.o3.bc	34.56	2.78	ind	0.08	0.02	0.05
elevator_spec9_product29_safe.cil.o3.bc	23.7	10.81	ind	0.14	0.02	0.22
elevator_spec1_product03_safe.cil.o0.bc	231.96	11.9	ind	0.39	0.02	0.22
elevator_spec1_product27_safe.cil.o0.bc	298.7	16.6	ind	0.48	0.02	0.26
elevator_spec1_product19_safe.cil.o3.bc	50.83	8.46	ind	0.07	0.01	0.06
elevator_spec1_product17_safe.cil.o3.bc	33.41	5.68	ind	0.07	0.01	0.06
elevator_spec1_product27_safe.cil.o3.bc	31.42	6.26	ind	0.07	0.01	0.06
elevator_spec1_product01_safe.cil.o3.bc	94	7.9	ind	0.08	0.01	0.06
elevator_spec1_product03_safe.cil.o3.bc	68.34	7.77	ind	0.08	0.01	0.06
elevator_spec1_product31_safe.cil.o3.bc	123.46	14.01	ind	0.12	0.01	0.1
elevator_spec1_product19_safe.cil.o0.bc	200.74	36.72	ind	0.29	0.01	0.23
elevator_spec1_product11_safe.cil.o0.bc	246.57	18.87	ind	0.33	0.01	0.24
elevator_spec1_product17_safe.cil.o0.bc	410.57	19.2	ind	0.4	0.01	0.2
elevator_spec1_product01_safe.cil.o0.bc	376.01	18.1	ind	0.44	0.01	0.18
elevator_spec1_product11_safe.cil.o3.bc	70.06	12.81	ind	0.09	0	0.06
elevator_spec1_product23_safe.cil.o3.bc	333.93	34.86	ind	0.13	0	0.09
elevator_spec1_product21_safe.cil.o3.bc	217.8	26.96	ind	0.2	0	0.09

Figure 4: Verification statistics of oV, OPTVERIFY and uV (part 3 of 3)