

Speed and Precision in Range Analysis

Victor Hugo Sperle Campos, Raphael Ernani Rodrigues,
Igor Rafael de Assis Costa and Fernando Magno Quintão Pereira

Department of Computer Science – UFMG – Brazil
{victorsc,raphael,igor,fernando}@dcc.ufmg.br

Abstract. Range analysis is a compiler technique that determines statically the lower and upper values that each integer variable from a target program may assume during this program’s execution. This type of inference is very important, because it enables several compiler optimizations, such as dead and redundant code elimination, bitwidth aware register allocation, and detection of program vulnerabilities. In this paper we empirically evaluate an inter-procedural, context-sensitive range analysis algorithm that we have implemented in the LLVM compiler. During the effort to produce an industrial-quality implementation of our algorithm, we had to face a constant tension between precision and speed. The foremost goal of this paper is to discuss the many engineering choices that, due to this tension, have shaped our implementation. Given the breath of our evaluation, we believe that this paper contains the most comprehensive empirical study of a range analysis algorithm ever presented in the compiler related literature.

1 Introduction

Range analysis is a compiler technique whose objective is to determine statically, for each program variable, limits for the minimum and maximum values that this variable might assume during the program execution. Range analysis is important because it enables many compiler optimizations. Among these optimizations, the most well-known are dead and redundant code elimination. Examples of redundant code elimination include the removal of array bounds checks [3, 13, 27] and overflow checks [22]. Additionally, range analysis is also used in bitwidth aware register allocation [1, 19, 26], branch prediction [18] and synthesis of hardware for specific applications [4, 12, 14, 23]. Because of this importance, the programming language community has put much effort in the design and implementation of efficient and precise range analysis algorithms.

However, the compiler related literature does not contain a comprehensive evaluation of range analysis algorithms that scale up to entire programs. Many works on this subject are limited to very small programs [14, 21, 23], or, given their theoretic perspective, have never been implemented in production compilers [9, 10, 24, 25]. There are implementations of range analysis that deal with very large programs [2, 6, 13, 16]; nevertheless, because these papers focus on applications of range analysis, and not on its implementation, they do not provide a

thorough discussion about their engineering decisions. A noticeable exception is the recent work of Oh *et al.* [17], which discusses a range analysis algorithm developed for C programs that can handle very large benchmarks. Oh *et al.* present an evaluation of the speed and memory consumption of their implementation. In this paper we claim to push this discussion considerably further.

We have implemented an industrial-quality range analysis algorithm in the LLVM compiler [11]. While designing and implementing our algorithm we had to face several important engineering choices. Many approaches that we have used in an attempt to increase the precision of our implementation would result in runtime slowdowns. Although we cannot determine the optimum spot in this design space, given the vast number of possibilities, we discuss our most important implementation decisions in Section 3. Section 3.1 shows how we can improve runtime and precision substantially by processing data-flow information in the strongly connected components that underly our constraint system. Section 3.2 discuss the importance of choosing a suitable intermediate representation when implementing a sparse data-flow framework. Section 3.3 compares the intra-procedural and the inter-procedural versions of our algorithm. The role of context sensitiveness is discussed in Section 3.4. Finally, Section 3.5 discusses the different widening strategies that we have experimented with.

This work concludes a two years long effort to produce a solid and scalable implementation of range analysis. Our first endeavor to implement such an algorithm was based on Su and Wagner’s constraint system [24, 25]. However, although we could use their formulation to handle a subset of C-like constructs, their description of how to deal with loops was not very explicit. Thus, in order to solve loops we adopted Gawlitza *et al.*’s [9] approach. This technique uses the Bellman-Ford algorithm to detect increasing or decreasing cycles in the constraint system, and then saturates these cycles via a simple widening operator. A detailed description of our implementation has been published by Couto and Pereira [8]. Nevertheless, the inability to handle comparisons between variables, and the cubic complexity of the Bellman-Ford method eventually led us to seek alternative solutions to range analysis. This quest reached a pinnacle in the present work, which we summarize in this paper.

2 Brief Description of our Range Analysis Algorithm

The Interval Lattice. Following Gawlitza *et al.*’s notation, we shall be performing arithmetic operations over the lattice $\mathcal{Z} = \mathbb{Z} \cup \{-\infty, +\infty\}$, where the ordering is naturally given by $-\infty < \dots -1 < 0 < 1 < \dots < +\infty$. We let meet and join be the min and max operators respectively. For any $x > -\infty$ we define:

$$\begin{aligned} x + \infty &= \infty & x - \infty &= -\infty \\ x \times \infty &= \infty \text{ if } x > 0 & x \times \infty &= -\infty \text{ if } x < 0 \\ 0 \times \infty &= 0 & (-\infty) \times \infty &= \text{not defined} \end{aligned}$$

From the lattice \mathcal{Z} we define the product lattice \mathcal{Z}^2 , partially ordered by the subset relation \sqsubseteq , and defined as $\mathcal{Z}^2 = \emptyset \cup \{[z_1, z_2] \mid z_1, z_2 \in \mathcal{Z}, z_1 \leq z_2, -\infty < z_2\}$. The objective of range analysis is to determine a mapping $I : V \mapsto \mathcal{Z}^2$ from




Fig. 1. Our implementation of range analysis. Rounded boxes are optional modules. The grey box is a module implemented in LLVM; the other parts are our contributions.

the set of integer program variables V to intervals, such that, for any variable $v \in V$, if $I(v) = [l, u]$, then, during the execution of the target program, any value i assigned to v is such that $l \leq i \leq u$.

A Holistic View of our Range Analysis Algorithm. Figure 1 gives a global view of our range analysis algorithm. We perform range analysis in a number of steps, some of which are optional. The optional parts improve the precision of the range analysis, at the expense of a longer running time. In Section 3 we discuss in more detail these tradeoffs.

We will illustrate the mandatory parts of the algorithm via the example program in Figure 2. More details about each phase of the algorithm will be introduced in Section 3, when we discuss our engineering decisions. Figure 2(a) shows an example program taken from the partition function of the quicksort algorithm used by Bodik *et al.* [3]. Figure 2(b) shows one possible way to represent this program internally. As we explain in Section 3.2, a good program representation helps range analysis to find more precise results. In this example we chose a program representation called Extended Static Single Assignment form [3], which lets us solve range analysis via a path sensitive algorithm. This representation uses the ϕ -functions typical in SSA form programs [7], plus *future*s (**ft**), which we shall define later. Figure 2(c) shows the constraints that we extract from the intermediate representation seen in part (b) of this figure. From these constraints we build the *constraint graph* in Figure 2(d). This graph is the main data-structure that we use to solve range analysis. For each variable v in the constraint system, the constraint graph has a node n_v . Similarly, for each constraint $v = f(\dots, u, \dots)$ in the constraint system, the graph has an *operation node* n_f . For each constraint $v = f(\dots, u, \dots)$ we add two edges to the graph: $\overrightarrow{n_u n_f}$ and $\overrightarrow{n_f n_v}$. Some edges in the constraint graph are dashed. These are called *control dependence edges*. If a constraint $v = f(\dots, \mathbf{ft}(u), \dots)$ uses a *future* bound from a variable u , then we add to the constraint graph a control dependence edge $\overrightarrow{n_u n_f}$. The final solution to this instance of the range analysis problem is given in Figure 2(e).

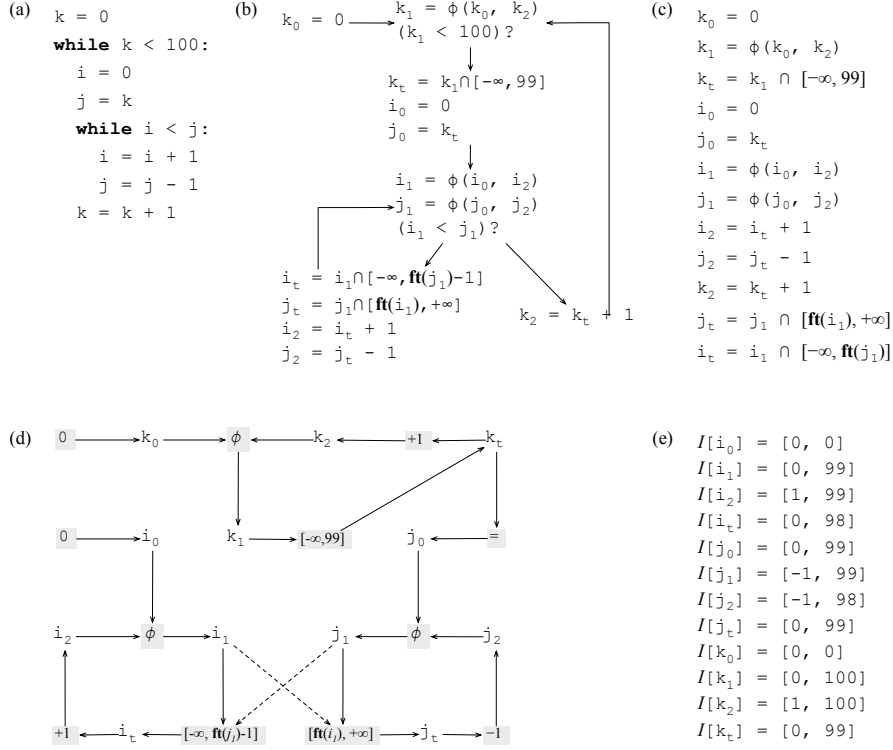


Fig. 2. Range analysis by example. (a) Input program. (b) Internal compiler representation. (c) Constraints of the range analysis problem. (d) The constraint graph. (e) The final solution.

The Micro Algorithm. We find the solution given in Figure 2(e) in a process that we call the micro algorithm. This process is divided into three sub-steps: (i) growth analysis; (ii) future resolution and (iii) narrowing analysis.

Growth analysis. The objective of growth analysis is to determine the growth behavior of each program variable. There are four possible behaviors: (a) the variable is bound to a constant interval, such as k_0 in Figure 2(b). (b) The variable is bound to a decreasing interval, i.e., an interval whose lower bound decreases. This is the case of j_1 in our example. (c) The variable is bound to an increasing interval, i.e., its upper bound increases. This is the case of i_1 in the example. (d) The variable is bound to an interval that expands in both directions. The growth analysis uses an infinite lattice, i.e., \mathcal{Z}^2 . Thus, a careless implementation of an algorithm that infers growth patterns might not terminate. In order to ensure termination, we must rely on a technique called *widening*, first introduced by

$$\frac{Y = X \sqcap [l, \mathbf{ft}(V) + c] \quad I(V)_{\uparrow} = u}{Y = X \sqcap [l, u + c]} \quad u, c \in \mathbb{Z} \cup \{-\infty, +\infty\}$$

$$\frac{Y = X \sqcap [\mathbf{ft}(V) + c, u] \quad I(V)_{\downarrow} = l}{Y = X \sqcap [l + c, u]} \quad l, c \in \mathbb{Z} \cup \{-\infty, +\infty\}$$

Fig. 3. Rules to replace futures by actual bounds. Given an interval $\iota = [l, u]$, we let $\iota_{\downarrow} = l$, and $\iota_{\uparrow} = u$

$$\frac{I(V)_{\downarrow} = -\infty \quad e(V)_{\downarrow} > -\infty}{I(V) \leftarrow [e(V)_{\downarrow}, I(V)_{\uparrow}]} \quad \frac{I(V)_{\downarrow} > e(V)_{\downarrow}}{I(V) \leftarrow [e(V)_{\downarrow}, I(V)_{\uparrow}]}$$

$$\frac{I(V)_{\uparrow} = +\infty \quad e(V)_{\uparrow} < +\infty}{I(V) \leftarrow [I(V)_{\downarrow}, e(V)_{\uparrow}]} \quad \frac{I(V)_{\uparrow} < e(V)_{\uparrow}}{I(V) \leftarrow [I(V)_{\downarrow}, e(V)_{\uparrow}]}$$

Fig. 4. Cousot and Cousot’s narrowing operator. Function $e(V)$ is an abstract evaluation, on the interval lattice, of the instruction that produces V .

Cousot and Cousot as a key component of abstract interpretation [5]. There are many different widening strategies. We discuss some of them in Section 3.5.

Future resolution. In order to learn information from comparisons between variables, such as $i < j$ in Figure 2(a), we bind some intervals to *futures*. Futures are symbolic limits, which will be replaced by actual numbers once we finish the growth analysis. The ranges found by the growth analysis tells us which variables have fixed bounds, independent on the intersections in the constraint system. Thus, we can use actual limits to replace intersections bounded by futures. Figure 3 shows the rules to perform these substitutions. In order to correctly replace a future $\mathbf{ft}(v)$ that limits a variable v' , we need to have already applied the growth analysis onto v . Had we considered only data dependence edges, then it would be possible that v' be analyzed before v . However, because of control dependence edges, this case cannot happen. The control dependence edges ensure that any topological ordering of the constraint graph either places N_v before $N_{v'}$, or places these nodes in the same strongly connected component. For instance, in Figure 2(d), variables j_1 and i_t are in the same SCC only because of the control dependence edges.

Narrowing analysis. The growth analysis associates very conservative bounds to each variable. Thus, the last step of our algorithm consists in narrowing these intervals. We accomplish this step via Cousot and Cousot’s classic narrowing operator [5, p.248], which we show in Figure 4.

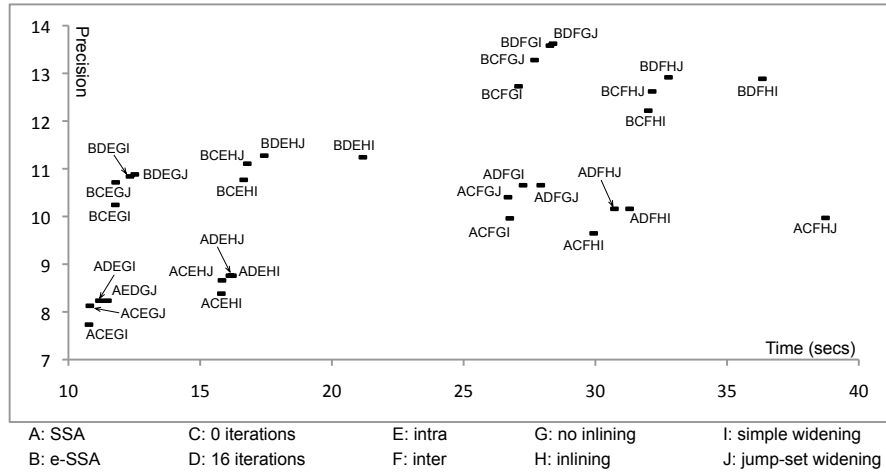


Fig. 5. Design space exploration: precision (percentage of bitwidth reduction) versus speed (secs) for different configurations of our algorithm analyzing the SPEC CPU 2006 integer benchmarks.

3 Design Space

As we see from a cursory glance at Figure 1, our range analysis algorithm has many optional modules. These modules give the user the chance to choose between more precise results, or a faster analysis. Given the number of options, the design space of a range analysis algorithm is vast. In this section we try to cover some of the most important tradeoffs. All the numbers that we show have been obtained as the average of 15 runs in an Intel Core 2 Quad processor with 2.4 GHz, and 3.5 GB of main memory. Figure 5 plots, for the integer programs in the SPEC CPU 2006 benchmark suite, precision versus speed for different configurations of our implementation. Our initial goal when developing this analysis was to support a bitwidth-aware register allocator. Thus, we measure precision by the average number of bits that our analysis allows us to save per program variable. It is very important to notice that we do not consider constants in our statistics of precision. In other words, we only measure bitwidth reduction in variables that a constant propagation step could not remove.

3.1 Strongly Connected Components

The greatest source of improvement in our implementation is the use of strongly connected components. To propagate ranges across the constraint graph, we fragment it into strongly connected components, collapse each of these components into single nodes, and sort the resulting directed acyclic graph topologically. We




Fig. 6. (Left) Bars give time to run our analysis without building strong components divided by time to run the analysis on strongly connected components. (Right) Bars give precision, in bitwidth reduction, that we obtain with strong components, divided by the precision that we obtain without them.

then solve the range analysis problem for each component individually. Once we have solved a component, we propagate its ranges to the next components, and repeat the process until we walk over the entire constraint graph. It is well-known that this technique is essential to speedup constraint solving algorithms [15, Sec 6.3]. In our case, the results are dramatic, mostly in terms of speed, but also in terms of precision. Figure 6 shows the speedup that we gain by using strong components. We show results for the integer programs in the SPEC CPU 2006 benchmark suite. In `xalancbmk`, the analysis on strong components is 450x faster.

According to Figure 6, in some cases, as in `bzip2`, strong components increase our precision by 40%. The gains in precision happen because, by completely resolving a component, we are able to propagate constant intervals to the next components, instead of propagating intervals that can grow in both directions. An example, in Figure ?? we pass the range $[0, 99]$ from variable k to the component that contains variable j . Had we run the analysis in the entire constraint graph, by the time we applied the growth analysis on j we would still find k bound to $[0, +\infty]$.

3.2 The choice of a program representation

If strong components account for the largest gains in speed, the choice of a suitable program representation is responsible for the largest gains in precision. However, here we no longer have a win-win condition: a more expressive program representation decreases our speed, because it increases the size of the target program. We have tried our analysis in two different program representations: the Static Single Assignment (SSA) form [7], and the Extended Static Single Assignment (e-SSA) form [3]. The SSA form gives us a faster, albeit more imprecise, analysis. Any program in e-SSA form has also the SSA core property: any variable name has at most one definition site. The contrary is not true: SSA form programs do not have the core e-SSA property: any use site of a variable that appears in a conditional test post-dominates its definition. The program in




Fig. 7. (Left) Bars give the time to run analysis on e-SSA form programs divided by the time to run analysis on SSA form programs. (Right) Bars give the size of the e-SSA form program, in number of assembly instructions, divided by the size of the SSA form program.

Figure 2(b) is in e-SSA form. The live ranges of variables i_1 and j_1 have been split right after the conditional test via the assertions that creates variables i_t and j_t . The e-SSA format serves well analyses that extract information from definition sites and conditional tests, and propagate this information forwardly. Examples include, in addition to range analysis, tainted flow analysis [20] and array bounds checks elimination [3].

Figure 7 compares these two program representations in terms of runtime. As we see in Figure 7(Left), the e-SSA form slows down our analysis. In some cases, as in `xalancbmk`, this slowdown increases execution time by 71%. Runtime increases for two reasons. Firstly, the e-SSA form programs are larger than the SSA form programs, as we show in Figure 7(Right). However, this growth is small: we did not verify any growth larger than 9% in any integer program of SPEC CPU 2006. Second, the e-SSA form program has futures; hence requiring the future resolution phase of our algorithm, which is not necessary in SSA form programs. Nevertheless, whereas the e-SSA form slows down the analysis runtime, its gains in precision are remarkable, as seen in Figure 8. These gains happen because the e-SSA format lets the analysis to use the results of comparisons to narrow the ranges of variables.

3.3 Intra versus Inter-procedural Analysis

A naive implementation of range analysis would be intra-procedural; that is, would solve the range analysis problem once per each function. However, we can gain in precision by performing it inter-procedurally. An inter-procedural implementation allows the results found for a function f to flow into other functions that f calls. Figure 9 illustrates the inter-procedural analysis for the program seen in Figure 2(a). The trivial way to produce an inter-procedural implementation is to insert into the constraint system assignments from the actual parameter names to the formal parameter names. In our example of Figure 9, our constraint




Fig. 8. The impact of the e-SSA transformation on precision for three different benchmark suites. Bars give the ratio of precision (in bitwidth reduction), obtained with e-SSA form conversion divided by precision without e-SSA form conversion.




Fig. 9. Example where an inter-procedural analysis is more precise than an intra-procedural analysis.

graph contains a flow of information from 0, the actual parameter, to k_0 , the formal parameter of function `foo`.

Figure 11 compares the precision of the intra and inter-procedural analyses for the five largest programs in three different categories of benchmarks: SPEC CPU 2006, the Stanford Suite ¹ and Bitwise [23]. Our results for the SPEC programs were disappointing: on average, for the five largest programs, the intra-procedural version of our analysis saves 5.23% of bits per variable. The inter-procedural version increases this number to 8.89%. A manual inspection of the SPEC programs reveals that this result is expected: these programs use many external library functions, which we cannot analyze, and their source codes do not provide enough explicit constants to power our analysis up. However, with numerical benchmarks we fare much better. On average, our inter-procedural al-

¹ <http://classes.engineering.wustl.edu/cse465/docs/BCCEexamples/stanford.c>




Fig. 10. Example where a context-sensitive implementation improves the results of range analysis.

gorithm reduces the bitwidth of the Stanford benchmarks by 36.24%. For Bitwise we obtain a bitwidth reduction of 12.27%. However, this average is lowered by two outliers: `edge_detect` and `sha`, which cannot be reduced. The Bitwise benchmarks were implemented by Stephenson *et al.* [23] to validate their bitwidth analysis. Our results are on par with those found by the original authors. The Bitwise programs contain only the `main` function; thus, different versions of our algorithm find the same results when applied onto these programs.

3.4 Context Sensitive versus Context Insensitive Analysis

Another way to increase the precision of range analysis is via a context-sensitive implementation. Context-sensitiveness allows us to distinguish different calling sites of the same function. Figure 10 shows why the ability to make this distinction is important for precision. In Figure 10(a) we have two different calls of function `foo`. An usual way to perform a data-flow analysis inter-procedurally is to create assignments between formal and actual parameters, as we show in Figure 10(b). If a function is called more than once, then its formal parameters will receive information from many actual parameters. We use the SSA’s ϕ -functions to bind this information together into a single flow. However, in this case the multiple assignment of values to parameters makes the ranges of these parameters very large, whereas in reality they are not. As an example, in Figure 10(b), variable k ends up associated with the range $[0, 10^5]$, but in reality this variable contains an interval that is only 100 units long. A way to circumvent this source of imprecision is via function inlining, as we show in Figure 10(c). The results that we can derive for the transformed program are more precise, as each input parameter is assigned a single value.

Figure 11 shows how function inlining modifies the precision of our results. It is difficult to find an adequate way to compare the precision of our analysis with, and without inlining. This difficulty stems from the fact that this transformation tends to change the target program too much. In absolute numbers, we




Fig. 11. The impact of whole program analysis on precision. Each bar gives precision in %bitwidth reduction.




Fig. 12. Runtime comparison between intra, inter and inter+inline versions of our algorithm. The bars are normalized to the time to run the intra-procedural analysis.

always reduce the bitwidth of more variables after function inlining. However, proportionally function inlining leads to a smaller percentage of bitwidth reduction for many benchmarks. In the Stanford Collection, for instance, where most of the functions are called in only one location, inlining leads to worse precision results. On the other hand, for the SPEC programs, inlining, even in terms of percentage of reduction, tends to increase our measure of precision.

Intra vs Inter-procedural runtimes. Figure 12(Right) compares three different execution modes. Bars are normalized to the time to run the intra-procedural analysis without inlining. On average, the intra-procedural mode is 28.92% faster than the inter-procedural mode. If we perform function inlining, then this difference is 45.87%. These numbers are close because our runtime is bound to the size of the strong components. We have observed that function inlining does not increase too much these components.

3.5 Choosing a Widening Strategy

We have implemented the widening operator used in the growth analysis in two different ways. The first way, which we call *simple*, is based on Cousot and Cousot's original widening operator [5]. This operator is shown in Figure 13,




Fig. 13. (Left) The lattice used in the simple widening strategy. (Right) Cousot and Cousot’s widening operator. We evaluate the rules from left-to-right, top-to-bottom, and stop upon finding a pattern matching.

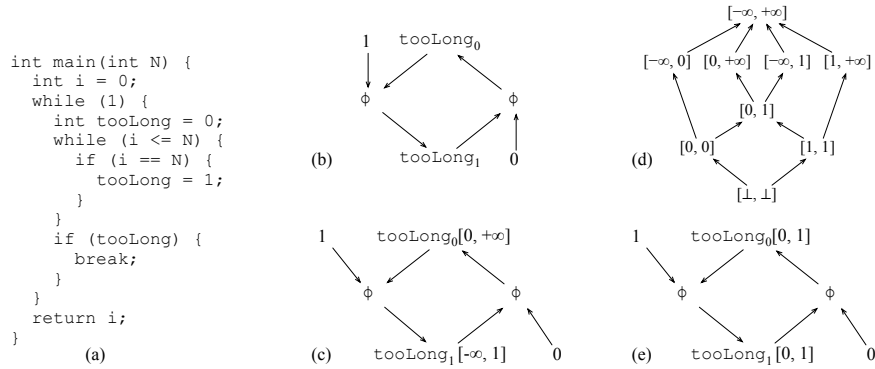


Fig. 14. An example where jump-set widening is more precise.

and it is the one used in Figure ??(b). The second widening strategy, which we call *jump-set widening* consists in using the constants that appear in the program text, in sorted order, as the next limits of each interval after widening is applied. This operator is common in implementations of range analysis [15, p.228]. There are situations in which jump-set widening produces better results than the simple operator. Figure 14 shows an example taken from the code of SPEC CPU **bzip2**. Part of the constraint graph of the program in Figure 14(a) is given in Figure 14(b). The result of applying the simple operator is shown in Figure 14(c). Jump-set widening would use the lattice in Figure 14(d), instead of the lattice in Figure 13(Right). This lattice yields the result given in Figure 14(e), which is more precise.

Another way to improve the precision of growth analysis is to perform a few rounds of abstract interpretation on the constraint graph, and to apply widening only if this process does not reach a fixed point. Each round of abstract

Benchmark	Size	0 + Simple	16 + Simple	0 + Jump	16 + Jump
hmmmer	38,409	9.98	11.40 (12.45)	10.98 (9.11)	11.40 (12.45)
gobmk	84,846	8.15	9.93 (17.92)	9.02 (9.64)	10.13 (19.54)
h264ref	97,494	12.58	13.11 (4.04)	13.00 (3.23)	13.11 (4.04)
xalancbmk	352,423	7.71	7.98 (3.38)	7.95 (3.02)	7.98 (3.38)
gcc	449,442	16.09	16.63 (3.25)	16.41 (1.95)	16.64 (3.31)

Fig. 15. Variation in the precision of our analysis given the widening strategy. The size of each benchmark is given in number of variable nodes in the constraint graph. Precision is given in percentage of bitwidth reduction. Numbers in parenthesis are percentage of gain over 0 + Simple.

interpretation consists in evaluating all the constraints, and then updating the intervals that change from one evaluation to the other. For instance, in Figure 14 one round of abstract interpretation, coupled with the simple widening operator, would be enough to reach the fixed point of that constraint system. We have experimented with 0 and 16 iterations before doing widening, and the overall result, for the programs in the SPEC CPU 2006 suite is given in Figure 5. Figure 15 shows some of these results in more detail for the five largest benchmarks in this collection. In general jump-set widening improves the precision of our results in non-trivial ways. Nevertheless, the simple widening operator preceded by 16 rounds of abstract interpretation in general is more precise than jump-set widening without any cycle of pre-evaluation, as we see in Figure 15.

4 Final Remarks

This paper presents what we believe is the most comprehensive evaluation of range analysis in the literature. Altogether we have experimented with 32 different configurations of our range analysis algorithm. Our implementation is publicly available at <http://code.google.com/p/range-analysis/>. This repository contains instructions about how to deploy and use our implementation. We provide a gallery of examples, including source codes, CFGs and constraint graphs that we produce for meaningful programs at <http://code.google.com/p/range-analysis/wiki/gallery>.

References

1. Rajkishore Barik, Christian Grothoff, Rahul Gupta, Vinayaka Pandit, and Raghavendra Udupa. Optimal bitwise register allocation using integer linear programming. In *LCPC*, volume 4382 of *Lecture Notes in Computer Science*, pages 267–282. Springer, 2006.
2. J. Bertrane, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, and X. Rival. Static analysis and verification of aerospace software by abstract interpretation. In *I@A*, pages 1–38. AIAA, 2010.

3. Rastislav Bodik, Rajiv Gupta, and Vivek Sarkar. ABCD: eliminating array bounds checks on demand. In *PLDI*, pages 321–333. ACM, 2000.
4. Jason Cong, Yiping Fan, Guoling Han, Yizhou Lin, Junjuan Xu, Zhiru Zhang, and Xu Cheng. Bitwidth-aware scheduling and binding in high-level synthesis. *Design Automation Conference, 2005. Proceedings of the ASP-DAC 2005. Asia and South Pacific*, 2:856–861, 18–21 Jan. 2005.
5. P. Cousot and R. Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, pages 238–252. ACM, 1977.
6. Patrick Cousot, Radhia Cousot, Jérôme Feret, Laurent Mauborgne, Antoine Miné, and Xavier Rival. Why does astrée scale up? *Form. Methods Syst. Des.*, 35(3):229–264, 2009.
7. Ron Cytron, Jeanne Ferrante, Barry K. Rosen, Mark N. Wegman, and F. Kenneth Zadeck. Efficiently computing static single assignment form and the control dependence graph. *TOPLAS*, 13(4):451–490, 1991.
8. Douglas do Couto Teixeira and Fernando Magno Quintao Pereira. The design and implementation of a non-iterative range analysis algorithm on a production compiler. In *SBLP*, pages 45–59. SBC, 2011.
9. Thomas Gawlitza, Jerome Leroux, Jan Reineke, Helmut Seidl, Gregoire Sutre, and Reinhard Wilhelm. Polynomial precise interval analysis revisited. *Efficient Algorithms*, 1:422 – 437, 2009.
10. Lies Lakhdar-Chaouch, Bertrand Jeannet, and Alain Girault. Widening with thresholds for programs with complex control graphs. In *ATVA*, pages 492–502. Springer-Verlag, 2011.
11. Chris Lattner and Vikram S. Adve. LLVM: A compilation framework for lifelong program analysis & transformation. In *CGO*, pages 75–88. IEEE, 2004.
12. G Lhairech-Lebreton, P Coussy, D. Heller, and E. Martin. Bitwidth-aware high-level synthesis for designing low-power dsp applications. In *ICECS*, pages 531–534. IEEE, 2010.
13. Fancesco Logozzo and Manuel Fahndrich. Pentagons: a weakly relational abstract domain for the efficient validation of array accesses. In *SAC*, pages 184–188. ACM, 2008.
14. S. Mahlke, R. Ravindran, M. Schlansker, R. Schreiber, and T. Sherwood. Bitwidth cognizant architecture synthesis of custom hardware accelerators. *Computer-Aided Design of Integrated Circuits and Systems*, 20(11):1355–1371, 2001.
15. Flemming Nielson, Hanne R. Nielson, and Chris Hankin. *Principles of Program Analysis*. Springer, 1999.
16. Hakjoo Oh, Lucas Brutschy, and Kwangkeun Yi. Access analysis-based tight localization of abstract memories. In *VMCAI*, pages 356–370. Springer, 2011.
17. Hakjoo Oh, Kihong Heo, Wonchan Lee, Woosuk Lee, and Kwangkeun Yi. Design and implementation of sparse global analyses for C-like languages. In *PLDI*, pages 229–238. ACM, 2012.
18. Jason R. C. Patterson. Accurate static branch prediction by value range propagation. In *PLDI*, pages 67–78. ACM, 1995.
19. Fernando Magno Quintao Pereira and Jens Palsberg. Register allocation by puzzle solving. In *PLDI*, pages 216–226. ACM, 2008.
20. Andrei Alves Rimsa, Marcelo D’Amorim, and Fernando M. Q. Pereira. Tainted flow analysis on e-SSA-form programs. In *CC*, pages 124–143. Springer, 2011.
21. Axel Simon. *Value-Range Analysis of C Programs: Towards Proving the Absence of Buffer Overflow Vulnerabilities*. Springer, 1th edition, 2008.

22. Marcos Rodrigo Sol Souza, Christophe Guillon, Fernando Magno Quintao Pereira, and Mariza Andrade da Silva Bigonha. Dynamic elimination of overflow tests in a trace compiler. In *CC*, pages 2–21, 2011.
23. Mark Stephenson, Jonathan Babb, and Saman Amarasinghe. Bitwidth analysis with application to silicon compilation. In *PLDI*, pages 108–120. ACM, 2000.
24. Zhendong Su and David Wagner. A class of polynomially solvable range constraints for interval analysis without widenings and narrowings. In *TACAS*, pages 280–295, 2004.
25. Zhendong Su and David Wagner. A class of polynomially solvable range constraints for interval analysis without widenings. *Theoretical Computer Science*, 345(1):122–138, 2005.
26. Sriraman Tallam and Rajiv Gupta. Bitwidth aware global register allocation. In *POPL*, pages 85–96. ACM, 2003.
27. Arnaud Venet and Guillaume Brat. Precise and efficient static array bound checking for large embedded c programs. *SIGPLAN Not.*, 39:231–242, 2004.