

Improving WCET Evaluation using Linear Relation Analysis *

Pascal Raymond, Claire Maiza, Catherine Parent-Vigouroux,
Erwan Jahier, Nicolas Halbwachs, Fabienne Carrier, Mihail Asavoae,
Rémy Boutonnet

Univ. Grenoble Alpes and CNRS, VERIMAG, Grenoble, France
first_name.last_name@univ-grenoble-alpes.fr

Abstract

The precision of a worst case execution time (WCET) evaluation tool on a given program is highly dependent on how the tool is able to detect and discard semantically infeasible executions of the program. In this paper, we propose to use the classical abstract interpretation-based method of *linear relation analysis* to discover and exploit relations between execution paths. For this purpose,

we add auxiliary variables (counters) to the program to trace its execution paths. The results are easily incorporated in the classical workflow of a WCET evaluator, when the evaluator is based on the popular *implicit path enumeration technique*. We use existing tools — a WCET evaluator and a linear relation analyzer — to build and experiment a prototype implementation of this idea.

2012 ACM Subject Classification Real-time systems software

Keywords and phrases Worst Case Execution Time estimation, Infeasible Execution Paths, Abstract Interpretation

Digital Object Identifier 10.4230/LITES.xxx.yyy.p

Received Date of submission. **Accepted** Date of acceptance. **Published** Date of publishing.

Editor LITES section area editor

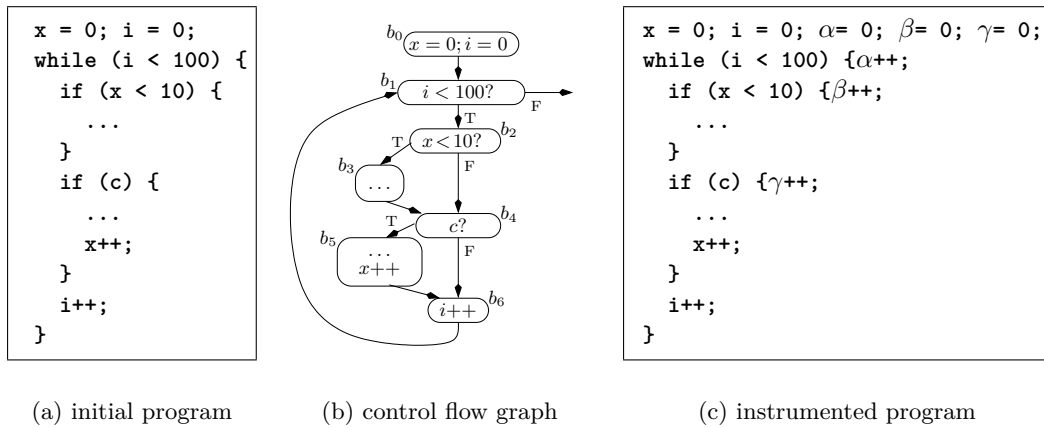
1 Introduction

The computation of a precise and safe approximation of the worst case execution time (WCET) of programs on a given architecture is an important step in the design of hard real-time systems [41]. It is part of the validation of the design, and a prerequisite for tasks scheduling. In this computation, over-approximation is mainly due to pessimistic abstraction of (1) complex hardware mechanisms (caches, pipeline) and (2) the program semantics (loop bounds, infeasible executions). Taking into account the target execution platform is, by far, the most difficult problem. It has been largely studied in the literature and remarkable tools exist, both in the academia [5, 27, 29] and in the industry [40].

In this paper, we specifically address the problem of taking into account the program semantics. The objective is to extract semantic properties that make some executions infeasible, and to exploit these properties in an existing WCET evaluator. It is generally admitted that such properties are easier to analyze on high-level code — e.g., C programs — than on binary, even if semantic analysis of executable code has been explored [3, 4, 36]. WCET evaluation is performed on object code in order to be able to take into account the execution architecture. This raises the problem of traceability between the source and the object code.

* This work is supported by the French research fundation (ANR) as part of the W-SEPT project (ANR-12-INSE-0001)





■ **Figure 1** Instrumenting an example program with counters

17 The most popular approach to evaluate the WCET is called *implicit path enumeration technique*
 18 (IPET) [28]. A micro-architectural analysis provides an evaluation of the duration of each basic
 19 block of the object-code control-flow graph. The WCET is then expressed as the solution of
 20 an integer linear programming problem (ILP) where the variables are the number of times each
 21 basic block is traversed during an execution. Relations between these variables come from the
 22 control-flow graph (flow equations) and from semantic “*flow facts*”, including at least *loop bounds*.
 23 Indeed, each loop in the program should have a constant bound to guarantee that the execution
 24 time is finite; such bounds may be provided by the user, or discovered by program analysis.

25 Hence, the IPET-based evaluation takes into account semantic properties expressed as linear
 26 constraints on counters. A natural idea is then to combine it with a semantic analysis devoted to
 27 the discovery of invariant linear relations. Polyhedra-based abstract interpretation [2, 8, 17, 20],
 28 also called *linear relation analysis* (LRA), is such an analysis. It is able to associate with each
 29 control point of a sequential program a system of linear inequalities (whose set of solutions is a
 30 convex polyhedron) satisfied by the numerical variables at this control point in any execution of
 31 the program.

32 Our proposal consists in applying LRA to a copy of the source program enriched with counter
 33 variables, and translate the obtained relations between counters into constraints to be added to
 34 the ILP. Let us illustrate this proposal with a simple example.

35 1.1 An example

36 Consider the program fragment of Figure 1.a with its control-flow graph (Fig. 1.b). Let’s add
 37 counters α , β , γ to the main basic blocks as done in the instrumented program Figure 1.c. These
 38 counters are initialized to 0 and incremented in their corresponding block. An LRA analysis of
 39 this instrumented program automatically discovers that the following relations are satisfied at the
 40 end of the program:

$$41 \quad \alpha = \mathbf{i} = 100, \gamma = \mathbf{x}, \beta + \gamma \leq 110, \gamma \geq 0, \beta \geq 0$$

42 The inequality $\alpha = 100$ gives the exact bound of the loop. More interestingly, $\beta + \gamma \leq 110$ means
 43 that there are at most 10 iterations of the loop where both blocks b_3 and b_5 are executed.

44 Assume the object code has the same control structure as the C program, i.e., the basic
 45 blocks of their control flow graphs are in an one-to-one correspondence. The standard WCET
 46 evaluation computes pessimistic execution times (say $t_i, i = 0..6$) of the basic blocks ($b_i, i = 0..6$),

47 and constructs the following ILP, where n_i (resp., $e_{i,j}$) denotes the number of occurrences of the
 48 basic block b_i (resp., the edge from b_i to b_j) in an execution of the program:

$$49 \quad w_{cet} = \max \sum_{i=0}^6 n_i \cdot t_i \quad , \quad \text{with the constraints} \quad \left\{ \begin{array}{ll} n_0 = 1 & , \quad e_{0,1} = n_0 \\ n_1 = e_{0,1} + e_{6,1} & , \quad e_{1,2} + 1 = n_1 \\ n_2 = e_{1,2} & , \quad e_{2,3} + e_{2,4} = n_2 \\ n_3 = e_{2,3} & , \quad e_{3,4} = n_3 \\ n_4 = e_{2,4} + e_{3,4} & , \quad e_{4,5} + e_{4,6} = n_4 \\ n_5 = e_{4,5} & , \quad e_{5,6} = n_5 \\ n_6 = e_{4,6} + e_{5,6} & , \quad e_{6,1} = n_6 \end{array} \right.$$

50

51 If we are able to maintain the correspondence between basic blocks in the source and the object
 52 code, i.e., to associate our counters α, β, γ with the variables of the ILP (n_2, n_3, n_5 respectively),
 53 we can add to the ILP the corresponding constraints: $n_2 = 100, n_3 + n_5 \leq 110$, which is likely to
 54 reduce the maximum value of the objective function¹.

55 1.2 Contents of the paper

56 In Section 2, we focus on some available tools, and experiment their semantics awareness on some
 57 simple examples. Two recent papers were dedicated to the state of the art related to semantic
 58 analyses for WCET estimation and infeasible path detection [1, 10]. Section 2.3 presents some
 59 more recent publications.

60 Our proposal consists in combining existing techniques, namely IPET-based WCET analysis
 61 and Linear Relation Analysis, recalled in Section 3, together with the specific tools that we used
 62 in our implementation. In Section 4, we explain how the counters are added and related to ILP
 63 counters thanks to debugging information provided by the compiler. Our implementation of
 64 the method is used to validate the approach on two existing benchmarks. We also investigated
 65 the robustness of the approach in presence of compiler optimizations. These experiments are
 66 summarized in Section 5. We conclude with some future works.

67 2 Existing tools

68 We have experimented with some existing tools, to evaluate their ability to discover and exploit
 69 semantic properties. Four tools have been considered, all of which go through similar steps:

- 70 1. extracting a control-flow graph from the object code,
- 71 2. performing a set of micro-architectural analyses to obtain execution times for each basic blocks,
- 72 3. using IPET to compute a safe WCET.

73 We compare these tools with respect to their capabilities to extract semantic properties to cut
 74 infeasible paths.

75 2.1 The tools

76 2.1.1 The Chronos Timing Analyzer

77 Chronos [27] is an academic tool developed at National University of Singapore. It takes as input
 78 a C program, performs limited data-flow analysis at C source code level to determine loop bounds,

¹ In fact, for this simple example, the results can be computed symbolically: concerning the standard evaluation, if the number of iterations in the loop ($= 100$) is given as a flow fact, the result will be $t_0 + 101t_1 + 100(t_2 + t_3 + t_4 + t_5 + t_6)$. Taking the additional constraint into account, we get $t_0 + 101t_1 + 100(t_2 + t_4 + t_6) + 100 \max(t_3, t_5) + 10 \min(t_3, t_5)$ thus improving the previous result by $90 \min(t_3, t_5)$.

79 and requests the user to provide this information when it fails. The semantic analysis in Chronos
 80 uses a pattern-based method to detect infeasible paths [39]. The so-called two-phase technique
 81 addresses infeasibility from a *conflicting pairs* point of view. In the first phase, an analysis detects
 82 some conflicts that capture the fact that two branches can not be taken along the same path. In
 83 the second phase, each conflicting pair relation is encoded into an ILP constraint.

84 **2.1.2 The Swedish Timing Analyzer**

85 SWEET² [29] is a research toolbox developed at Mälardalen Real-Time Research Center (MRTC).
 86 The main objective of SWEET is flow analysis, which computes flow-facts, i.e., information about
 87 loop bounds and infeasible paths in the program. The main technique to discover flow-facts is
 88 abstract execution [16]. Abstract execution is a form of context-sensitive abstract interpretation,
 89 because it uses a symbolic execution to produce context information for each loop iteration and
 90 function call. Instead of using the fixpoint engine of abstract interpretation, abstract execution
 91 executes the program in the abstract domain, merging the execution paths at certain points in
 92 the program. SWEET does not support LRA. It currently implements only the abstract domain
 93 of intervals.

94 **2.1.3 AbsInt - The aiT Tool**

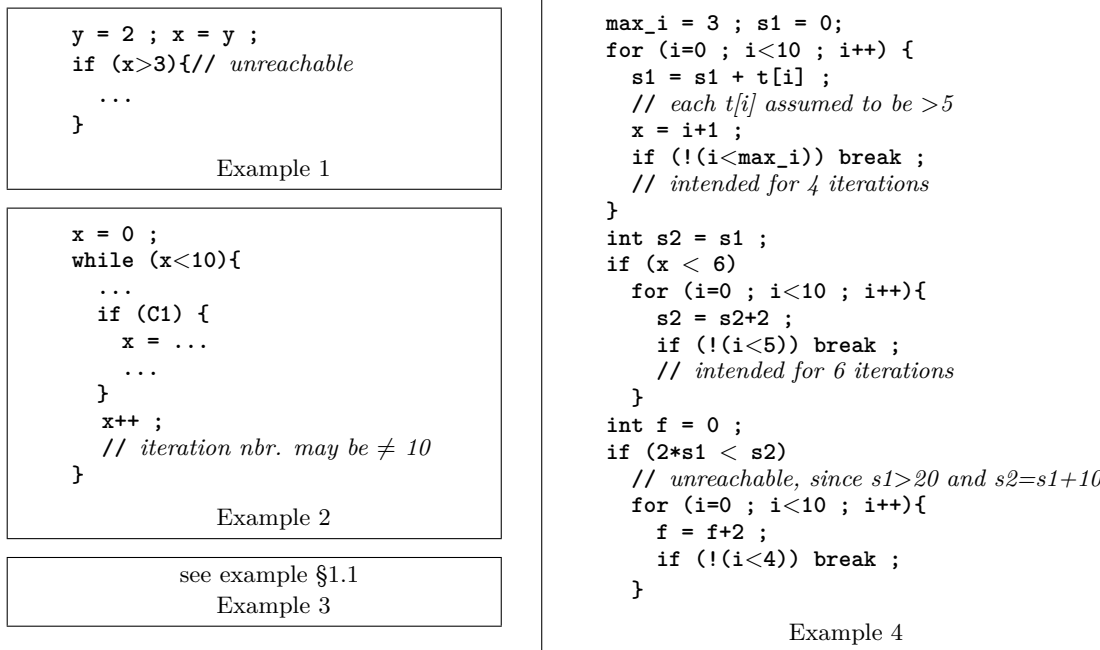
95 Developed by AbsInt³, aiT is the main industrial product for WCET analysis. It consists of a set
 96 of binary executables analyzers, which take the intrinsic cache and pipeline behavior into account.
 97 Concerning semantic analysis, aiT uses a value analysis based on intervals [13] to compute safe
 98 ranges of values for the program variables. aiT uses this information to determine loop bounds
 99 and detect infeasible paths. The approach towards computing loop bounds is not general, but it
 100 handles loop patterns. In order to gain precision, aiT pre-processes each loop by transforming
 101 its body into a function, in order to expose the iteration contexts. The key element in this
 102 transformation is to identify the loop index and to set it as a function parameter. Then, an
 103 interval analysis computes the ranges for all the loop variables. The loop transformation is based
 104 on loop patterns, which depends on the particularities of the architecture (e.g., parameter order)
 105 or on the loop structure (e.g., for-loops, triangular-loops, branch conditions). aiT is able to detect
 106 infeasible paths using the results of the value analysis, like conditions made infeasible because of
 107 the computed intervals.

108 **2.1.4 oRange, the flow fact analyzer of OTAWA**

109 OTAWA [5] is an academic toolbox, developed at IRIT (University of Toulouse), designed as
 110 a generic framework to develop static analyses for WCET computation. Although OTAWA
 111 implements several approaches to WCET computation, the one based on IPET is the most mature.
 112 OTAWA relies on an auxiliary tool, called oRange [9], to compute loop bounds. oRange analyses
 113 C code. As a first phase, oRange detects loop indices and constructs a normal form: a symbolic
 114 expression of the bound independently of the call context. In a second phase, by an abstract
 115 execution, a syntactic tree is built in function of a full or partial call context. It combines loop
 116 bounds and conditional expressions as numeric or symbolic expressions. Finally, the tree is
 117 computed in the full context in order to produce a file in the specific flow-facts format FFX [42].

² www.mrtc.mdh.se/projects/wcet

³ AbsInt GmbH www.absint.com/ait/



■ **Figure 2** Various cases of semantic infeasibilities

2.2 Some experiments

In order to evaluate the capabilities of these tools to detect infeasible paths, we have applied each of them to programs containing various situations of semantic infeasibility. These situations are given on Figure 2:

- Example 1 is a case where simple pattern-based method may fail, since constant propagation is needed.
- Example 2 may be a problem for pattern-based methods for finding iteration numbers, since the apparent index x is modified.
- Example 3 is our introductory example of §1.1.
- Example 4 is a fragment of code generated by the SCADÉ⁴ compiler, from a design manipulating arrays. On one hand, the loops are exited from inside, which makes complex the evaluation of iteration numbers. On the other hand, the third loop is unreachable because of some non-trivial arithmetic conditions.

Table 1 summarizes the results of the tools on these examples. On Example 1, Chronos is unable to detect dead code⁵. Example 2 is correctly analyzed only by SWEET, because it unrolls loops. None of the tools is able to find the property of Example 3. On Example 4, Chronos requires manual annotations for loop bounds; oRange estimates that both loops are iterated 10 times; SWEET and aiT find the exact loop bounds; only aiT detects dead code.

In this paper, we propose a method and a tool-chain that is able to discover the infeasible paths of these 4 examples, namely, infeasible paths that depend on a semantic analysis and that may concern distant program points.

⁴ www.esterel-technologies.com/products/scade-suite

⁵ We used the available version of Chronos. Some additional work has been done that complement the infeasible path analysis [6, 37], which is not part of the available version.

	Chronos	SWEET	oRange	aiT
Example 1	-	✓	✓	✓
Example 2	-	✓	-	-
Example 3	-	-	-	-
Example 4				
nbr. 1st loop	-	4	10	4
nbr. 2nd loop	-	5	10	5
dead code	-	-	-	✓

■ **Table 1** Results of tools on programs of Figure 2

139 2.3 Other approaches

140 An extended state of the art related to semantic analyses for WCET estimation can be found
 141 in [1] and a general survey of infeasible path detection in [10]. We complement them with some
 142 more recent publications.

143 Several recent works make use of SMT solvers [23, 37]. The idea is to ask the solver if
 144 the worst-case path obtained by the ILP solver is feasible. Whenever the path is infeasible, a
 145 corresponding constraint is added to the ILP. As in our approach, adding constraints does not
 146 always mean that the WCET is refined (2 paths may have the same WCET). In [18], the whole
 147 path analysis is done through SMT solving instead of ILP: infeasible path analysis and worst-case
 148 path analysis are merged in one step. Path execution time is expressed as an SMT problem, and
 149 the question asked is no longer “is this path feasible?”, but “is there a feasible path longer than
 150 K ?”, where K is a given constant (which is adjusted, e.g., by binary search). In [32, 35], a similar
 151 approach is taken, by asking this question to a bounded model checker.

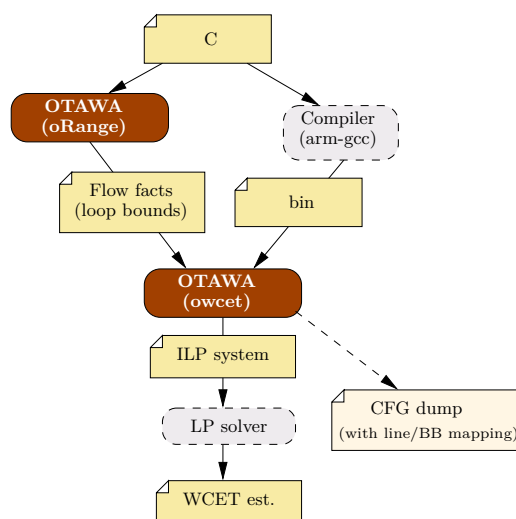
152 3 Used techniques and tools

153 This section presents the existing techniques and tools used in our prototype: OTAWA implements
 154 the classical IPET-based WCET evaluation, and PAGAI performs Linear Relation Analysis.

155 3.1 WCET evaluation with OTAWA

156 The WCET estimation work-flow (Figure 3) involves a compiler, a Linear Program solver, and
 157 two tools from the OTAWA toolbox: oRange and owcet.

- 158 ■ **Compilation:** the source C code is compiled by a third party tool; for this experiment, we use
 159 a cross compiler from the GNU Compiler Collection (arm-elf-gcc 4.4.2), but other compiler
 160 can be used, provided that it produces ELF code (Executable and Linkable Format), with
 161 debugging information in DWARF format.
- 162 ■ **oRange** is a *data flow analyzing tool*, dedicated to the discovery of *loop bounds*. Bounds are
 163 stored in the OTAWA *flow facts* format (FFX).
- 164 ■ **owcet** is the OTAWA command dedicated to the WCET evaluation. The main steps of this
 165 tool, not detailed in Figure 3, are:
 - 166 ■ the construction of the control-flow graph (CFG) of the object code; during the construction,
 167 and thanks to debugging information, basic blocks (BB) are associated (if possible) to lines
 168 in the source program; thanks to this correspondence, the loop bounds computed by oRange
 169 are translated into control flow constraints in the CFG. The annotated CFG can be dumped
 170 in a file, allowing other tools to exploit it.



■ **Figure 3** Ottawa WCET estimation work-flow

- 171 ■ the micro-architectural analysis, which associates a local WCET estimation to each BB of
- 172 the CFG.
- 173 ■ the construction of the Integer Linear Programming (ILP) system; as in the introduction
- 174 example (§1.1) the resulting system gathers (1) structural constraints (CFG structure), (2)
- 175 loop bounds constraints (from oRange flow facts) (3) the objective function to be maximized
- 176 (sum of BB counters weighted by their local WCET).
- 177 ■ ILP solver: the ILP system is then solved by a third-party tool; OTAWA integrates and uses
- 178 LP_SOLVE⁶ (any other equivalent tool can be used).

179 3.2 Linear Relation Analysis with PAGAI

180 3.2.1 Principles of LRA

181 Linear Relation Analysis [8] is a classical program analysis, based on abstract interpretation [7]. It
 182 is able to discover, at each control point of a sequential program, a conjunction of linear relations
 183 (equalities and inequalities) invariantly satisfied by the numerical variables at this point. Classical
 184 algorithms are used to propagate linear systems over the statements of the program. Several
 185 causes may result in information lost:

- 186 ■ the analysis safely ignores non-linear expressions in assignments and tests;
- 187 ■ the analysis performs a *convex hull* at control path junctions, instead of propagating the
- 188 disjunction of incoming information. It means that the propagated value is the most precise
- 189 conjunction of linear relations implied by both incoming systems;
- 190 ■ to avoid infinite propagation along loops, the classical *widening-narrowing* method is applied
- 191 to guess a safe approximation of the limit. Note that, unlike in symbolic execution [23] or
- 192 SMT methods [18], loops are not unrolled.

⁶ web.mit.edu/lpsolve/doc

193 3.2.2 Applying LRA to our example

194 We do not detail further the techniques applied in LRA, and refer the reader to the bibliography.
 195 We just show the main steps of the analysis of our example of Figure 1. Let's consider the control
 196 point at the entry of the while loop. The first step of the analysis straightforwardly computes the
 197 first iterate:

$$198 \quad x = i = \alpha = \beta = \gamma = 0$$

199 Its propagation through the loop body provides, with a convex hull at the end of the conditional:

$$200 \quad i = \alpha = \beta = 1, x = \gamma, 0 \leq \gamma \leq 1$$

201 Now a convex hull with the first iterate gives the second iterate at the entry of the loop:

$$202 \quad i = \alpha = \beta, 0 \leq \alpha \leq 1, x = \gamma, 0 \leq \gamma \leq \alpha$$

203 Instead of continuing the iterations, a first widening/narrowing step is performed (using “lookahead
 204 widening” [14]), which provides:

$$205 \quad i = \alpha = \beta, x = \gamma, 0 \leq \gamma \leq \alpha \leq 100, \gamma \leq 10$$

206 Now, the “else” branch of the test $x < 10$ becomes feasible, and a second widening/narrowing step
 207 is performed, providing:

$$208 \quad i = \alpha, x = \gamma, \beta + \gamma \leq \alpha + 10, \beta \leq \alpha, \gamma \leq \alpha \leq 100$$

209 which is found invariant after one more propagation. Propagated to the end of the program, it
 210 becomes:

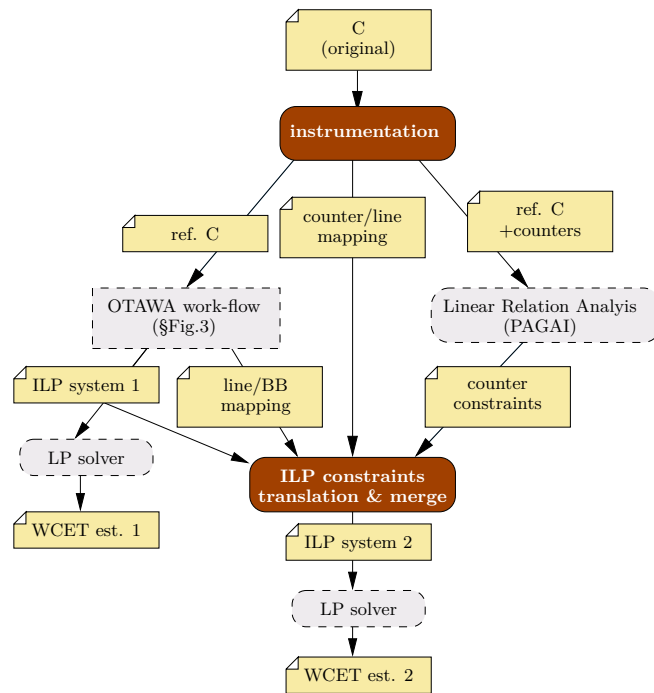
$$211 \quad i = \alpha = 100, x = \gamma \leq 100, \beta + \gamma \leq 110$$

212 3.2.3 LRA and loop bounds

213 It may happen, like in the previous example, that LRA discovers a bound to a loop counter, thus
 214 providing an essential information for WCET evaluation. However, finding loop bounds is *not* our
 215 main goal in this work, as the method is intrinsically unable to discover *non linear* relations, which
 216 drastically limits its capability to find loop bounds. As a matter of fact, in presence of nested loops,
 217 the number of executions of the body of the innermost loop is not linear in the constants of the
 218 program. For instance, in the program fragment “`for (i=0; i<n; i++){for (j=i; j<n; j++){...}}`”
 219 the number of executions of the body of the innermost loop is $n(n+1)/2$, which cannot be found
 220 by LRA.

221 The LRA method must then be used together with some other method able to bound nested
 222 loops. We can use existing tool, such that oRange that comes with OTAWA, or more basically
 223 user-given bounds, given as pragmas in the code.

224 There exist also approaches based on polyhedra manipulation to find loop bounds, such as
 225 the one proposed in [38, 30]: it consists in building a polyhedral upper approximation P of the
 226 iteration domain, i.e., the set of possible valuations of loop counters (in the previous example,
 227 $P = \{(i, j) \mid 0 \leq i \leq j \leq n - 1\}$). Under realistic assumptions concerning the determinism
 228 of the program, the number of executions of the innermost loop is bounded by the number of
 229 integer points in P , and algorithms are available to compute this number. Notice that LRA can
 230 be combined with this approach, since it can discover linear invariants reducing the iteration
 231 domain, thus improving the precision of the result. Notice also that LRA can deal with parameters
 232 (symbolic bounds, like n in our example), an issue specifically addressed by [38].



■ **Figure 4** Instrumentation and analysis workflow

3.2.4 The PAGAI prototype analyzer

Several tools performing LRA are available ([2, 12, 19, 20] to cite a few). Here, we use the PAGAI prototype analyzer, which implements the basic LRA together with recent improvements like “lookahead widening” [14] and SMT-based “path focusing” [19]. PAGAI analyses LLVM code [24] produced from a C program (thanks to Clang⁷), and is able to return discovered properties at the C level. PAGAI may be used with other abstract domains than general linear systems — like octagons [33] — thanks to the common interface APRON [21].

4 Adding and tracing counters

4.1 The proposed workflow

Figure 4 shows the proposed workflow for the experiment. It involves two existing components: timing analysis with OTAWA (left) and program analysis with PAGAI (right). Two new tools have been developed to complete the workflow: a front-end (top, *Instrumentation*), which produces the input for the analyzers (OTAWA and PAGAI), and a back-end (*ILP translation & merge*), which gathers the results into a more constrained ILP system, and obtains a possibly enhanced WCET estimation.

These tools are detailed in this section. We illustrate the successive steps of the method by detailing the processing of an example program, called `lcdnum.c`, extracted from TacleBench programs suite [15]. The main program is given in Figure 5. It calls a function `num_to_lcd`, the execution time of which is taken into account by OTAWA.

⁷ <http://clang.llvm.org/>

252 **4.2 Instrumented program version**

253 The goal of the front-end (“instrumentation”, Figure 4, top) is to produce, from the original C
 254 code, a reference C program. Some semantics preserving transformations of the source code are
 255 necessary or advisable, in order to use properly the analyzers, and trace the information between
 256 them.

- 257 ■ Some transformations are purely lexical, and do not change the program structure: because
 258 the standard ELF/DWARF traceability mechanism is line-based, line breaks are introduced to
 259 isolate each atomic statements on its own line.
- 260 ■ Some transformations that modify the control structure are necessary because of the limitation
 261 of the analyzers. For instance, a single-return statement per function is mandatory for exploiting
 262 the results of PAGAI: this unique control point is the place where counters invariants actually
 263 express properties on the whole execution of the function. Other transformations are required
 264 because of the limitation of both OTAWA and PAGAI: the control structure (CFG) must
 265 be statically known, which forbids dynamic computation of program pointers. In particular,
 266 “switch/case” statements must be rewritten into a static control structure based on “if” and
 267 “goto” statements.
- 268 ■ Another transformation is desirable in our case: the current version of PAGAI does not handle
 269 inter-procedural analysis. In order to exploit the plain capacity of this tool to find invariants,
 270 a light-weight solution is to inline function calls at the source level. This transformation is
 271 indeed hardly admissible in real-life, but it must be seen here as a “trick” to reach our goal
 272 (study the ability of LRA to detect infeasible executions).

273 The front-end produces the reference C code in two flavors.

- 274 ■ The reference C code with counters (Figure 4, right) is instrumented with auxiliary counters,
 275 in the same manner as in the introduction example (§ 1.1). The present version introduces a
 276 counter for each sequential block in the program control flow. However, some strategy could
 277 be used to reduce the number of counters by targeting blocks that are more likely to have an
 278 influence [43].
- 279 ■ The reference C code without counters (Figure 4, left) is the same code, where all lines
 280 related the counters (declaration, initialization and incrementation) have been commented
 281 out. This method ensures a semantic equivalence between the programs analyzed by OTAWA
 282 and PAGAI: since they only differ on the side-effect-free local variables, these programs are
 283 naturally input/output equivalent. Moreover, at least at the source level, the two programs
 284 are also structurally equivalent: a block in the reference C code is executed if and only if the
 285 corresponding block (marked with a counter c) is executed in the reference C program with
 286 counters. This property becomes false in general at the binary level, since the C compiler may
 287 modify the control structure: this well-know problem of traceability is discussed later.
- 288 ■ An auxiliary file is generated, that contains the mapping between each counter and its
 289 corresponding source line in the reference C code.

| *Example:* Applied to our example program (Figure 5), our instrumentation front-end calls the

C preprocessor, eliminates the multiple returns and switches (only within `num_to_lcd`, not shown), and produces the reference C programs. The first one (without counters) is shown on Figure 6; the second (not shown) is exactly the same with uncommented lines involving counters. An auxiliary file (not shown) simply lists the pairs “counter/line” (e.g., (`cptr_main_1,144`), (`cptr_main_2, 147`)).

The first version is provided to OTAWA. Loop bounds computation by `oRange` is optional, which allows us to check if PAGAI is able to find them on its own. OTAWA calls the gcc compiler (here with `-O0` optimization level), builds the CFG of the object code, performs the micro-architectural analysis, and builds the ILP problem.

PAGAI is applied to the second version of the program, and returns the following invariants:

```
-10+cptr_main_2 = 0
-10+cptr_main_4 = 0
5-cptr_main_3 >= 0
```

The first equation finds the exact loop bound (which may also be found by `oRange`). The second equation is structural (from the shape of the source CFG, `cptr_main_2` and `cptr_main_4` are equal). The third property is new, and expresses, in particular, that the function `num_to_lcd` is called at most 5 times.

290 4.3 Tracing back the counters

291 The back-end (“ILP constraints translation & merge”, Figure 4, bottom) gathers the information
292 coming from OTAWA and PAGAI:

- 293 ■ Thanks to the counter/C-line mapping provided by the front-end, and the C-line/binary-
294 BB mapping provided by OTAWA (through the ELF/DWARF information), a counter/BB
295 mapping is built. Note that this mapping is partial, and deliberately pessimistic: depending
296 on the compilation process, it may happen that a counter is associated either to zero or to
297 several binary basic blocks. In this case, the counter is simply ignored: only counters that are
298 associated to one single BB are retained.

Example (cont.): Table 2 shows the mapping between counters and blocks that is built by our back-end.

- 299 ■ The linear constraints on the retained counters are then translated literally into linear con-
300 straints on BB, and added to the basic ILP system provided by OTAWA.

Example (cont.): The translation of the constraints discovered by PAGAI is the following:

```
x4_main = 10;
x6_main = 10;
x5_main <= 5;
```

- 301 ■ At last, both systems are solved and the corresponding estimations can be compared.

Example (cont.): After a second call to `LP_SOLVE`, the final result is printed:

```
Estimation WITHOUT PAGAI: 1540
Estimation WITH     PAGAI: 945
```

302 4.4 Traceability and optimization

303 In our framework, traceability is the ability to relate execution paths in the binary code (bin.
304 CFG) to execution paths in the source code (source CFG).

305 Some optimizations performed by the compiler may strongly modify the control structure and
306 thus alter traceability: loop unrolling, block replication, out-of-order execution. This is why most

```

unsigned char num_to_lcd( unsigned char a ) ;

volatile unsigned char IN = 120;
volatile unsigned char OUT;
int main( void ) {
    int i;
    unsigned char a;
    for( i=0; i< 10; i++ ) {
        a = IN;
        if(i<5) {
            a = a &0x0F;
            OUT = num_to_lcd(a);
        }
    }
    return 0;
}

```

■ **Figure 5** The initial lcdnum.c program

```

133 int main(void) {
134     int i ;
135     unsigned char a ;
136     unsigned char tmp ;
137     int __retres4 ;
138     //int cptr_main_1 = 0;
139     //int cptr_main_2 = 0;
140     //int cptr_main_3 = 0;
141     //int cptr_main_4 = 0;
142     //int cptr_main_5 = 0;
143     //cptr_main_1 ++; #line 144
144     i = 0;
145     while (i < 10) {
146         //cptr_main_2 ++; #line 147
147         a = (unsigned char )IN;
148         if (i < 5) {
149             //cptr_main_3 ++; #line 150
150             a = (unsigned char )((int )a & 15);
151             tmp = num_to_lcd(a);
152             OUT = (unsigned char volatile )tmp;
153         }
154         //cptr_main_4 ++; #line 155
155         i ++;
156     }
157     //cptr_main_5 ++; #158
158     __retres4 = 0;
159     return (__retres4);
160 }

```

■ **Figure 6** The reference lcdnum.c program

line number(s)	block(s)	reliable	counter
136,144	1	yes	cptr_main_1
145	1;2	no	
147;148	4	yes	cptr_main_2
150;151;152	5	yes	cptr_main_3
155	6	yes	cptr_main_4
158;159;160	3	yes	cptr_main_5

■ **Table 2** Mapping between counters and blocks

307 of the related works assume no compiler optimization to guarantee a perfect matching between
 308 the two CFGs.

309 However forbidding optimization is not satisfactory in real-time domains, where execution
 310 times have to be predictable, but also short. For a standard compiler like gcc, the observed
 311 speed-up between no optimization (-O0 option) and a standard level of optimization (-O1) is
 312 around two.

313 The most satisfactory solution would be a compiler that provides a precise traceability even in
 314 case of CFG optimization. Some work have been done to design and/or adapt the compilation
 315 process for this purpose, for instance [26, 31, 22].

316 Unfortunately, off-the-shelf standard compilers such as gcc hardly provide a precise and reliable
 317 information in case of CFG optimization. The idea is then to use the compiler options in order to
 318 forbid (as far as possible) CFG transformations, but still allow other optimizations, in particular
 319 all that concern data management.

320 The gcc compiler proposes numerous options to control optimizations, but there hardly exists
 321 a comprehensive and exhaustive description of their effects and inter-dependencies. For this
 322 experiment, we have empirically defined a *customized level* (called **C0** in the sequel). We started
 323 from the standard -O1 level, and remove about 20 individual optimizations using the `-fno` directive
 324 (see appendix A). We cannot guarantee that this customized level will preserve the CFG for all
 325 programs, but the method is safe: as explained in Section 4.3, a counter (and then a source code
 326 line) that is not associated to exactly one basic block of the binary code is simply ignored. As a
 327 consequence, the only risk is to loose information that would have made the WCET estimation
 328 tighter. Note that this statement suppose that the gcc debugging information is reliable, which is
 329 indeed unprovable, but *empirically* reasonable.

Example: When applying the **C0** method to our running example, we get a 100% traceability. As
 a consequence, the interesting counter property (`5-cptr_main_3 >= 0`) can still be translated
 into a BB constraint (`x11_main <= 5;`) leading to the final result:

```
Estimation WITHOUT PAGAI: 641
Estimation WITH     PAGAI: 421
```

On this example, we observe that code optimization leads to a initial WCET estimation 2.4x
 smaller (641 vs 1540). The traceability is preserved and the improvement due to the counter
 analysis is of the same order (34.3% vs 38.6%).

330 5 Experiments

331 5.1 Benchmarks

332 We tested our approach on programs from the TacleBench [11], a set of C programs widely used
 333 in the WCET community. A first check has been made to retain only purely sequential programs
 334 that compile “out of the box”: 53 applications of the 58 in the TacleBench⁸

335 For each program, we try to estimate the WCET of all functions appearing in the code,
 336 including the top-level one (main). For each function, inner function calls are recursively inlined at
 337 the C level (see Section 4.2). Recursive functions are rejected during this step, and not considered
 338 for WCET analysis.

⁸ The 5 missing applications are OS and/or architecture dependent.

339 Our goal is to study the influence of our counter-based method (Fig. 4) on a classical estima-
 340 tion (Fig. 3). A prerequisite is therefore that a *reference* estimation exists; hence the programs
 341 for which the basic WCET estimation fails are not selected. The OTAWA estimation may fails
 342 because of unsupported programming features (pointer arithmetics), or because the analysis does
 343 not terminates before a chosen timeout (2 hours).

344 After this initial selection, 589 functions (out of 639) from the 53 programs of the TacleBench
 345 suite are retained.

346 5.2 Experimental setup

347 The proposed framework as presented on Fig. 4 has numerous parameters (C code instrumenta-
 348 tion, linear analysis tuning, compiler optimization etc.) leading to a combinatorial numbers of
 349 possibilities. For this systematic experiment, we focus only on two kinds of parameters: those
 350 that influence the precision of linear analysis, and those that influence the traceability. The other
 351 parameters are fixed once and for all as follows:

352 **OTAWA hardware model:** our goal is not to bench or “stress” OTAWA in terms of hardware.

353 We only want it to give an initial IPET system in which we will insert flow facts discovered
 354 via LRA. In order to maximize the number of test benches for which OTAWA gives an initial
 355 ILP in reasonable time, we consider a very simple, cache-free, ARM-based architecture.

356 **Misc CFG transformations:** some CFG transformations are necessary, due to limitations of
 357 OTAWA (switch statements not supported) and /or PAGAI (multiple return statements). This
 358 transformations are performed using the CIL library [34].

359 **Inlining:** because the current version of PAGAI has limited support for inter-procedural analysis,
 360 function calls are systematically inlined. This transformation is also implemented using the
 361 CIL library. This method improves the precision of the analysis, but makes the analysis much
 362 costly in time and memory.

363 **Loop bounds:** as explained in 3.2.3, our method is intrinsically unable to bound nested loops,
 364 so a complementary method is necessary to find loop bounds. For this purpose, we can use
 365 oRange, but it appears that the CFG transformations performed using CIL strongly alters its
 366 performance⁹. In order to maximize the size of the benchmark we thus systematically exploit,
 367 when available, the user pragmas given in source code. Nevertheless, we made a complementary
 368 experiment, without using pragmas nor oRange, in order to identify the cases where LRA is
 369 sufficient to bound the execution time.

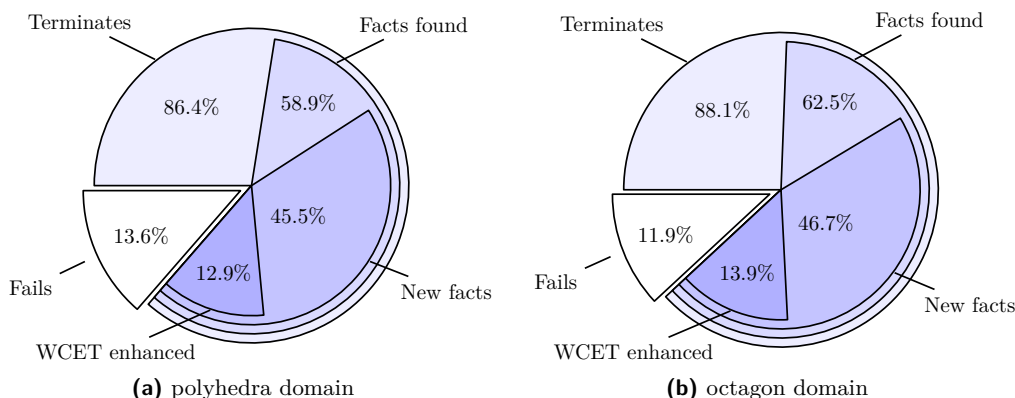
370 5.3 Lessons learnt

371 This section presents the lessons learnt form the experiment, by focusing on several points:
 372 the ability of the linear analysis to discover “flow facts”, and hopefully to enhance the WCET
 373 estimation; the influence of the abstract domain on the analysis; the ability of linear analysis to
 374 discover loop bounds, and finally the influence of compiler optimizations on traceability.

375 5.3.1 Linear analysis and flow facts discovery

376 When traceability allows it, the constraints discovered by linear analysis are directly translated
 377 into *flow facts* giving information on the (im)possible execution paths. These flow facts may be
 378 useless if they are redundant with the structural constraints, otherwise they are *new facts*, giving

⁹ The CIL tool normalizes the code by using only unbounded *while* and *break* statements, that are badly handled by oRange. However oRange performs well for the original programs, made of human-written *for* loops.



■ **Figure 7** LRA analysis statistics on 589 functions, for the two relational abstract domains.

■ **Table 3** Some WCET improvement results (full table page 25)

Ref	Initial WCET	Box			Octagons			Polyhedra		
		Δ	Imp ^t	Time	Δ	Imp ^t	Time	Δ	Imp ^t	Time
cr.2	227K	111K	48.7	<1s	111K	48.7	4s	111K	48.7	1s
md.13	2648	0	0.0	<1s	1920	72.5	<1s	1920	72.5	<1s
gs.9	6934	0	0.0	1s	738	10.6	4m	4428	63.8	29m
an.0	466M	0	0.0	4s	4M	0.8	7m	115M	24.6	1m
mp.9	52M	27M	51.1	56m	-	-	-	-	-	-
md.14	51K	0	0.0	2m	5K	10.4	21m	-	-	-
md.5	13M	0	0.0	2m	-	-	-	3M	21.0	35m


379 non trivial information on the execution paths. However, even new facts can be useless if they do
 380 not concern the worst case execution path. A utility has been developed to check whether the
 381 facts discovered by LRA analysis are new or not. Each fact is checked by adding its negation to
 382 the set of structural constraints: the fact is redundant if and only if the system becomes infeasible.

383 Figure 7 gives statistic on the behavior of the LRA method, for the two relational domains
 384 (octagon and polyhedra). Let's focus on the polyhedra case first (a). The PAGAI tool terminates
 385 for 509 cases out of 589 (86.4%); for the missing cases (13.6%), it runs out of resources in memory
 386 or time. Flow facts are found in 347 cases, and at least one fact is new for 268 ones; finally, new
 387 facts lead to a WCET improvement for 76 cases. Statistics are similar for the octagon domain,
 388 except that it terminates more often: this explains why the WCET is enhanced more often with
 389 octagons, even if this domain is less precise.

390 A possible conclusion is that LRA, when it works, is actually good at finding non redundant
 391 semantic facts (more than half of the time, when it terminates), but that those facts do not
 392 necessarily lead to a WCET improvement (about 15% of the termination cases).

393 5.3.2 Abstract domains

394 The main goal of the experiment is to observe the influence of the linear analysis on the WCET
 395 estimation. The linear analysis performed by PAGAI is parameterized by the choice of an abstract
 396 domain to represent the possible values of the counters. Two domains proposed by PAGAI are
 397 *relational*, and thus are likely to express relations between our counters and the original variables
 398 in the programs:



■ **Figure 8** WCET improvement and analysis time depending on abstract domains (b=box, o=octagons, p=polyhedra)

- 399 ■ The polyhedra domain is the most precise since it can handle any linear relation, and its
 400 algorithmic cost is exponential in the worst case.
 401 ■ The octagon domain handles intervals and bounded pairwise sums or differences. It is less
 402 precise, but has a polynomial cost in the worst case: $O(n^3)$ in time, and $O(n^2)$ in space.

403 To be exhaustive, we also consider the domain *box*, which handles only intervals. Since this
 404 domain is non-relational, it is intrinsically unable to relate our additional counters to the program
 405 variables. The flow facts that can be discovered with the box domain are thus limited (basically,
 406 counters stuck down to zero, which correspond to dead code).

407 The WCET estimation is improved by at least one domain for 90 functions. The gain ranges
 408 from negligible (0.1%) to interesting (around 10%) or even huge (more than 50%). We limit here
 409 the comments to the cases where the enhancement is greater 0.8%. The detailed results for these
 410 60 cases are given in appendix (table 6, page 25), and a selection of typical cases is given in table 3.

411 The experiment gives some interesting information:

- 412 ■ The interest of the *box domain* is very limited: it is an indirect way of performing constant
 413 propagation and dead code “pruning”. Most of the time it gives no improvement (42 out of 60,
 414 e.g., md.13, gs.9). However, since it is the cheapest domain, it may give results when other
 415 domains fail (6 times, e.g., mp.9).
 416 ■ When both *octagons* and *polyhedra* terminate, they often give the same result (34 out of 60
 417 cases, e.g., cr.2, md.13). However there are some cases (12 out of 60, e.g., gs.9), where the
 418 expressiveness of polyhedra is actually useful (constraint involving 3 or more variables, and
 419 pairwise relations with non unit coefficients).
 420 ■ In compliance with the theoretical complexity, *octagons* may terminate while *polyhedra* fails
 421 (7 cases, e.g., md.14). Nevertheless, there is also one case where *octagons* fail while *polyhedra*
 422 works (md.5). This is due to the fact that the cost of octagons is almost always cubic in the
 423 number of variables, while the exponential cost of polyhedra is rarely reached in practice.

■ **Table 4** Impact of compiler optimizations on WCET and LRA

Ref	O0			CO				
	Initial WCET	Best WCET	Best Imp ^t	Opt. speedup	Initial WCET	Best WCET	Best Imp ^t	Traceability
md.13	2648	728	72.5	3.3x	791	215	72.8	100% of 2
an.0	466M	351M	24.6	3.0x	157M	116M	25.9	100% of 44
cr.2	227K	116K	48.7	2.3x	97K	50K	48.7	41% of 24
md.5	13M	10M	21.0	3.4x	4M	3M	15.0	80% of 40
ex.2	278K	224K	19.2	1.3x	218K	218K	0.0	46% of 13

5.3.3 Loop bounds

LRA is intrinsically limited to the discovery of single loop bounds (cf. 3.2.3). We made a complementary experiment to check if and when LRA actually finds such loops. For this experiment, we only consider the short-list of programs from Table 6 where PAGAI terminates when using a relational domain (octagon and polyhedra); as a matter of fact, using the *box* domain is irrelevant since it can't find any loop bound other than 0.

For these 54 programs, we have:

- computed the *loop level*, which is maximal depth of nested loops appearing in the program (0: no loop at all, 1: only single loops, 2 or more: nested loops);
- launched our tool without using `oRange` nor user-pragmas. The LRA analysis is performed twice: with the octagon and the polyhedra domain, and we keep only the best result.

Table ?? (page 26) lists the results; the column “pagai” simply indicates if the analysis give a bounded WCET, since the WCET value is, in this case, the same as the one in Table 6.

There are 10 test cases that are loop-free, and thus with no bounds to found. There are 25 programs with only single loops (level=1); these are the cases where PAGAI is supposed to find bounds, and it actually does it for most of the cases (19 out of 25). In fact, PAGAI find the bounds for all loops that are semantically guarded by a counter condition, that is, *for* loops or equivalent. The cases where PAGAI does not find bounds are those where the loop is guarded by a *points-to* condition (e.g., `while (*p++)`).

We expected PAGAI to not bound any program with a loop level greater than 1, which is the case except for one program (ex.2). In fact this example is a “false counter-example”: the loop depth is *syntactically* 2, but the inner-loop appears in a branch which is never executed. The loop depth is then *semantically* 1.

5.3.4 Optimization level and traceability

The main focus of this work is the influence of linear analysis on the precision of the WCET. Nevertheless, since analysis is performed at the C level, the problem of the traceability between the C and the binary code must be considered. Forbidding any optimization is not an option in real-time domain. We argue that a well-chosen set of optimizations can lead to a reasonable compromise between traceability and program speed-up.

For all functions that give some enhancement on the non-optimized code, we run the experiments using the custom optimization (CO) level defined in 4.4. Since the counter analysis is completely independent to the compilation method, the linear relations found are the same, and the ability to enhance the WCET estimation is only due to traceability.

The detailed results of this experiment are given in appendix (table 8, page 27), and a selection of typical cases is given in table 4. The table gathers the results obtained with the non-optimized

459 binary code O0, and the optimized one CO. For each optimization level, the table gives:

- 460 ■ the *Initial WCET*, in cpu-cycles, computed by OTAWA,
- 461 ■ the *Best WCET*, enhanced thanks to the properties discovered with PAGAI, with some abstract
- 462 domain,
- 463 ■ the corresponding **Improvement** percentage.

464 The table also shows the *Optimization speedup*, which is the ratio between the initial O0 and the
 465 initial CO estimation, i.e., it measures the gain obtained just because of the compilation, before
 466 applying the counter method. Finally, for CO compilation, the table gives an information on the
 467 *Traceability*: the percentage of counters introduced for LRA at C level, that are actually associated
 468 to some basic block, at binary level. Traceability in the O0 mode is not shown in the table as it is
 469 always 100%.

470 The interesting information given by the experiment are:

- 471 ■ Even if the CO level is very limited (subset of O1 level, and a fortiori of O2), it generates a
- 472 fairly optimized code: the speedup is mostly between 2x and 4x.
- 473 ■ In most of the cases (53 out of 60) traceability is 100%, and one can observe an enhancement
- 474 due to LRA similar to the one obtained with O0 code. Indeed, this improvement is obtained
- 475 on the CO initial WCET, which is already much smaller than the O0 one (e.g., md.13, an.0).
- 476 ■ In some cases, traceability is partly lost, but remain sufficient to enhance the estimation (4
- 477 cases, e.g., cr.2, md.5).
- 478 ■ Finally, for 3 cases, partial traceability leads to no enhancement (e.g., ex.2).

479 **6 Conclusion and future work**

480 Linear Relation Analysis is a powerful technique to discover invariant linear relations between
 481 numerical variables of a program. On the other hand, the classical evaluation of WCET using
 482 Implicit Path Enumeration Technique is based on expressing the WCET as the solution of an
 483 Integer Linear Program, the variables of which are counters associated with the basic blocks of the
 484 program. So, the idea of adding these counters as auxiliary variables in the program, and using
 485 the results of LRA as semantic flow-facts to be added to the ILP, is rather natural. Our goal, in
 486 this paper, was to conduct a light-weight experiment — by combining existing tools — to evaluate
 487 the benefits of the approach. Secondly, such an experiment raised the question of traceability,
 488 since semantic flow-facts are discovered on the source program, while the WCET is evaluated on
 489 the executable code. The conclusion of this experiment on public benchmarks is manyfold:

- 490 ■ LRA finds new semantic facts in many examples (46%), but many of these new facts don't
- 491 influence the evaluated WCET. However, the WCET is improved on a significant subset (almost
- 492 14%) of the examples, and the improvement is often interesting.
- 493 ■ the traceability problems can be safely dealt with, using the debugging information provided
- 494 by the compiler; this is the case even in presence of strong compiling optimizations, as long as
- 495 these optimizations don't modify too much the control structure of the program.

496 This work could be continued in several directions.

- 497 ■ It would be interesting to limit the number of counters, as the cost of LRA can be exponential
- 498 in the number of variables. Of course, counters which are structurally related by flow equations
- 499 can be saved, but their cost is low in polyhedra computations (they are linked to each other by
- 500 equations). An appealing idea would be to introduce counters on the branches of a conditional,
- 501 only when these branches appear to have strongly different execution times, a measure that is
- 502 roughly available after the micro-architectural analysis [43].
- 503 ■ Existing LRA analyzers (like PAGAI) are generally not inter-procedural, which forced us to
- 504 inline the procedures in our experiments. An inter-procedural version of LRA must be studied

to solve this problem. The *relational* nature of LRA is surely an advantage, since a procedure can be associated a summary as an input-output relation. Summaries of called procedures can then be used in the caller, in a bottom-up fashion.

- Traceability is still a concern, which would benefit from a better cooperation of the compiler [25].

References

- 1 Mihail Asavaoe, Claire Maiza, and Pascal Raymond. Program semantics in model-based WCET analysis: A state of the art perspective. In *13th International Workshop on Worst-Case Execution Time Analysis, WCET 2013, July 9, 2013, Paris, France*, volume 30 of *OASICS*, pages 32–41. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2013.
- 2 Roberto Bagnara, Elisa Ricci, Enea Zaffanella, and Patricia M. Hill. Possibly not closed convex polyhedra and the Parma polyhedra library. In M. V. Hermenegildo and G. Puebla, editors, *9th International Symposium on Static Analysis, SAS'02*, Madrid, Spain, September 2002. LNCS 2477. doi: 10.1007/3-540-45789-5_17.
- 3 Gogul Balakrishnan and Thomas W. Reps. DIVINE: DIScovering variables IN executables. In *Verification, Model Checking, and Abstract Interpretation, VMCAI 2007*, pages 1–28, Nice, France, January 2007.
- 4 Gogul Balakrishnan, Thomas W. Reps, David Mel-ski, and Tim Teitelbaum. WYSINWYX: what you see is not what you execute. In *Verified Software: Theories, Tools, Experiments, VSTTE 2005*, pages 202–213, Zurich, Switzerland, October 2005.
- 5 Clément Ballabriga, Hugues Cassé, Christine Rochange, and Pascal Sainrat. OTAWA: An open toolbox for adaptive WCET analysis. In *SEUS*, 2010.
- 6 Duc-Hiep Chu, Joxan Jaffar, and Rasool Maghareh. Precise cache timing analysis via symbolic execution. In *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS)*, pages 1–12, 2016.
- 7 Patrick Cousot and Radia Cousot. Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *4th ACM Symposium on Principles of Programming Languages, POPL'77*, Los Angeles, January 1977.
- 8 Patrick Cousot and Nicolas Halbwachs. Automatic discovery of linear restraints among variables of a program. In *5th ACM Symposium on Principles of Programming Languages, POPL'78*, Tucson (Arizona), January 1978.
- 9 Marianne de Michiel, Armelle Bonenfant, Hugues Cassé, and Pascal Sainrat. Static loop bound analysis of C programs based on flow analysis and abstract interpretation. In *IEEE Int'l Conf. on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, 2008.
- 10 Sun Ding, Hee Beng Kuan Tan, and Kaiping Liu. A survey of infeasible path detection. In *Proceedings of the 7th International Conference on Evaluation of Novel Approaches to Software Engineering (ENASE 2012)*, Wroclaw, Poland, 29-30 June, 2012., pages 43–52, 2012.
- 11 Heiko Falk, Sebastian Altmeyer, Peter Hellinckx, Björn Lisper, Wolfgang Puffitsch, Christine Rochange, Martin Schoeberl, Rasmus Bo Sorensen, Peter Wagemann, and Simon Wegener. Taclebench: A benchmark collection to support worst-case execution time research. In *16th International Workshop on Worst-Case Execution Time Analysis, WCET 2016, July 5, 2016, Toulouse, France*, pages 2:1–2:10, 2016.
- 12 Paul Feautrier and Laure Gonnord. Accelerated invariant generation for C programs with Aspic and C2fsm. In *Tools for Automatic Program Analysis (TAPAS)*, Perpignan, France, September 2010.
- 13 Christian Ferdinand, Florian Martin, Christoph Cullmann, Marc Schlickling, Ingmar Stein, Stephan Thesing, and Reinhold Heckmann. New developments in WCET analysis. In *Program Analysis and Compilation*, pages 12–52, 2006.
- 14 Denis Gopan and Thomas Reps. Lookahead widening. In *CAV'06*, Seattle, 2006.
- 15 Jan Gustafsson, Adam Betts, Andreas Ermedahl, and Björn Lisper. The Mälardalen WCET benchmarks: Past, present and future. In *Proc. of WCET*, pages 136–146, 2010.
- 16 Jan Gustafsson, Andreas Ermedahl, Christer Sandberg, and Björn Lisper. Automatic derivation of loop bounds and infeasible paths for WCET analysis using abstract execution. In *RTSS*, 2006.
- 17 Nicolas Halbwachs, Yann-Eric Proy, and Patrick Roumanoff. Verification of real-time systems using linear relation analysis. *Formal Methods in System Design*, 11(2):157–185, August 1997.
- 18 Julien Henry, Mihail Asavaoe, David Monniaux, and Claire Maiza. How to compute worst-case execution time by optimization modulo theory and a clever encoding of program semantics. In *SIGPLAN/SIGBED Conference on Languages, Compilers and Tools for Embedded Systems 2014, LCTES '14*, pages 43–52, June 2014.
- 19 Julien Henry, David Monniaux, and Matthieu Moy. Pagai: A path sensitive static analyser. *Electr. Notes Theor. Comput. Sci.*, 289:15–25, 2012.
- 20 François Irigoien, Pierre Jouvelot, and Rémy Triolet. Semantical interprocedural parallelization: An overview of the PIPS project. In *ACM Int. Conf. on Supercomputing, ICS'91, Köln*, 1991.
- 21 Bertrand Jeannet and Antoine Miné. Apron: A library of numerical abstract domains for static analysis. In *Computer Aided Verification (CAV 2009)*, Grenoble, France, pages 661–667, June 2009.
- 22 Raimund Kirner, Peter Puschner, and Adrian Prantl. Transforming flow information during code optimization for timing analysis. *Journal on Real-Time Systems*, 45(1-2), 2010.
- 23 Jens Knoop, Laura Kovács, and Jakob Zwirchmayr. WCET squeezing: on-demand feasibility refinement for proven precise WCET-bounds. In

- Proceedings of the 21st International Conference on Real-Time Networks and Systems*, pages 161–170. ACM, 2013.
- 24 Chris Lattner and Vikram Adve. LLVM: a compilation framework for lifelong program analysis & transformation. In *CGO'04*, pages 75–86, Washington, DC, August 2004. IEEE Computer Society.
 - 25 Hanbing Li, Isabelle Puaut, and Erven Rohou. Traceability of flow information: Reconciling compiler optimizations and WCET estimation. In *22nd International Conference on Real-Time Networks and Systems, RTNS'14, Versailles, France, October 8-10, 2014*, 2014.
 - 26 Hanbing Li, Isabelle Puaut, and Erven Rohou. Tracing Flow Information for Tighter WCET Estimation: Application to Vectorization. In *21st IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*, page 10, Hong-Kong, China, August 2015. URL: <https://hal.inria.fr/hal-01177902>.
 - 27 Xianfeng Li, Liang Yun, Tulika Mitra, and Abhik Roychoudhury. Chronos: A timing analyzer for embedded software. *Sci. Comput. Program.*, 69(1-3):56–67, 2007.
 - 28 Yau-Tsun Steven Li and Sharad Malik. Performance analysis of embedded software using implicit path enumeration. *IEEE Trans. on Computer-Aided Design of Integrated Circuits and Systems*, 16(12), 1997.
 - 29 Björn Lisper. SWEET – a tool for WCET flow analysis. In *6th International Symposium On Leveraging Applications of Formal Methods, Verification and Validation (ISOLA)*, October 2014.
 - 30 Paul Lokuciejewski, Daniel Cordes, Heiko Falk, and Peter Marwedel. A fast and precise static loop analysis based on abstract interpretation, program slicing and polytope models. In *Proceedings of the CGO 2009, The Seventh International Symposium on Code Generation and Optimization*, pages 136–146, Seattle, Washington, USA, March 2009.
 - 31 Paul Lokuciejewski and Peter Marwedel. *Worst-Case Execution Time Aware Compilation Techniques for Real-Time Systems*. Springer, 2011. doi:10.1007/978-90-481-9929-7.
 - 32 Ravindra Metta, Martin Becker, Prasad Bokil, Samarjit Chakraborty, and R. Venkatesh. TIC: a scalable model checking based approach to WCET estimation. In *Proceedings of the 17th ACM SIGPLAN/SIGBED Conference on Languages, Compilers, Tools, and Theory for Embedded Systems, LCTES 2016, Santa Barbara, CA, USA, June 13 - 14, 2016*, pages 72–81, 2016. doi:10.1145/2907950.2907961.
 - 33 Antoine Miné. The octagon abstract domain. In *Proceedings of the Eighth Working Conference on Reverse Engineering, WCRE'01, Stuttgart, Germany, October 2-5, 2001*, page 310, 2001. doi:10.1109/WCRE.2001.957836.
 - 34 George C. Necula, Scott McPeak, Shree P. Rahul, and Westley Weimer. Cil: Intermediate language and tools for analysis and transformation of c programs. In R. Nigel Horspool, editor, *Compiler Construction*, pages 213–228, Berlin, Heidelberg, 2002. Springer Berlin Heidelberg.
 - 35 Pascal Raymond, Claire Maiza, Catherine Parent-Vigouroux, Fabienne Carrier, and Mihail Asavoae. Timing analysis enhancement for synchronous program. *Real-Time Systems*, pages 1–29, 2015.
 - 36 Jordy Ruiz and Hugues Cassé. Using SMT solving for the lookup of infeasible paths in binary programs (regular paper). In *Workshop on Worst-Case Execution Time Analysis, Lund, Sweden, 07/07/2015*, pages 95–104. OASiCs, Dagstuhl Publishing, July 2015.
 - 37 Thomas Sewell, Felix Kam, and Gernot Heiser. Complete, high-assurance determination of loop bounds and infeasible paths for WCET analysis. In *2016 IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS), Vienna, Austria, April 11-14, 2016*, pages 185–195, 2016. doi:10.1109/RTAS.2016.7461326.
 - 38 Björn Lisper Stefan Bygde, Andreas Ermedahl. An efficient algorithm for parametric WCET calculation. *Journal of Systems Architecture*, 57(6):614–624, May 2011.
 - 39 Vivy Suhendra, Tulika Mitra, Abhik Roychoudhury, and Ting Chen. Efficient detection and exploitation of infeasible paths for software timing analysis. In *DAC*, pages 358–363, 2006.
 - 40 Stephan Thesing, Jean Souyris, Reinhold Heckmann, Famantanantsoa Randimbivololona, Marc Langenbach, Reinhard Wilhelm, and Christian Ferdinand. An abstract interpretation-based timing validation of hard real-time avionics software. In *DSN*, pages 625–632, 2003.
 - 41 Reinhard Wilhelm, Jakob Engblom, Andreas Ermedahl, Niklas Holsti, Stephan Thesing, David Whalley, Guillem Bernat, Christian Ferdinand, Reinhold Heckmann, Tulika Mitra, Frank Mueller, Isabelle Puaut, Peter Puschner, Jan Staschulat, and Per Stenström. The worst-case execution-time problem - overview of methods and survey of tools. *ACM Trans. Embedded Comput. Syst. (TECS)*, 7(3), 2008.
 - 42 Jakob Zwirchmayr, Armelle Bonenfant, Marianne de Michiel, Hugues Cassé, Laura Kovács, and Jens Knoop. FFX: A portable WCET annotation language (regular paper). In *International Conference on Real-Time and Network Systems (RTNS), Pont-à-Mousson, 08/11/2012-09/11/2012*, pages 91–100, November 2012.
 - 43 Jakob Zwirchmayr, Pascal Sotin, Armelle Bonenfant, Denis Claraz, and Philippe Cuenot. Identifying Relevant Parameters to Improve WCET Analysis (regular paper). In *Workshop on Worst-Case Execution Time Analysis, Madrid, 08/07/2014*, pages 91–100. OASiCs, Dagstuhl Publishing, juillet 2014.

A Compiler optimization level

Options controlling optimizations are numerous and may vary a lot depending on the targeted processor and the compiler version. Options listed here are for the compiler used for our experiment (`arm-elf-gcc` (GCC) 4.4.2), with no guarantee that they apply directly to other compilers. Ensuring the coherence of a set of optimizations is technically hard, this is why we start with a predefined level of optimization (`-O1`) and remove optimizations that may modify the control structure (using the corresponding `-fno` flag). To select the options, we just rely on the user manual: we remove any optimization that mention a possible influence on the control structure or that may affect the precision of the debugging information (i.e. the association instruction/source line).

Loop transformations may drastically change the structure of the program, and are thus forbidden.

```
-fno-loop-block \  
-fno-loop-interchange \  
-fno-loop-strip-mine \  
-fno-move-loop-invariants \  
-fno-reschedule-modulo-scheduled-loops \  
-fno-unroll-loops \  
-fno-unroll-all-loops \  
-fno-unsafe-loop-optimizations \  

```

Miscellaneous CFG transformations concern dead code elimination, inlining, branch removal, block reordering etc.

```
-fno-dce \  
-fno-dse \  
-fno-guess-branch-probability \  
-fno-inline-small-functions \  
-fno-crossjumping \  
-fno-if-conversion \  
-fno-if-conversion2 \  
-fno-jump-tables \  
-fno-reorder-blocks \  
-fno-reorder-blocks-and-partition \  
-fno-unswitch-loops \  

```

SSA tree optimizations and misc. global optimizations have an indirect influence in the control structure, by removing, regrouping or re-ordering instructions. They also affect the precision of the debugging information (dwarf) which is the only information we have to relate the binary and the source code.

```
-fno-tree-builtin-call-dce \  
-fno-tree-ccp \  
-fno-tree-ch \  
-fno-tree-copyrename \  
-fno-tree-dce \  
-fno-tree-dominator-opts \  
-fno-tree-dse \  
-fno-tree-fre \  

```

```

554 -fno-tree-loop-distribution \
555 -fno-tree-loop-im \
556 -fno-tree-loop-ivcanon \
557 -fno-tree-loop-linear \
558 -fno-tree-loop-optimize \
559 -fno-tree-sra \
560 -fno-tree-ter \
561 -fno-auto-inc-dec \
562 -fno-cprop-registers \
563 -fno-defer-pop \
564 -fno-ipa-pure-const \
565 -fno-ipa-reference \
566 -fno-merge-constants \
567 -fno-split-wide-types \
568 -fno-unit-at-a-time \

```

569 **B** Experiment Results

570 Experiment was performed on 589 individual C functions extracted from the TACLeBench [11].
571 An improvement of the WCET estimation is observed for 90 functions (15% of the cases). This
572 section details the results for the 60 cases where the improvement is greater than 0.8%.

573 Table 5 contains label definitions to ease and shorten the reference to the bench functions:
574 the label (column 1), the source folder in the TACLeBench (column 2), and the function name
575 (column 3).

576 Table 6 contains the experiment results using the gcc `-O0` compilation level. The first column
577 holds the function label, and the second one holds the initial WCET estimation computed by
578 OTAWA. The remaining columns hold information related to the improvement obtained (or not)
579 with Linear relation analysis, using 3 different abstract domains: boxes (intervals), octagons
580 and polyhedra. For each domain, the table gives the improvement in number of cycles (Δ) and
581 percentage (Imp^t), and the time necessary to perform the LRA with PAGAI¹⁰. Numbers in bold
582 highlight the best improvements among various methods (box, octagons, polyhedra). Empty cells
583 ('-') mean that the corresponding case triggered the 2 hours timeout set for the experiment.

584 Table ?? gives information on the ability of PAGAI to discover loop bounds ; to obtain this
585 table, the experiments are re-played without the help of a external loop-bound method (neither
586 `oRange` nor the user-given pragmas). For each program, the table gives its loop level (maximal
587 depth of nested loops) and indicated wheter PAGAI finds a bounded WCET or not.

588 Finally, table 8 aims at observing the impact of compiler optimization on WCET estimation in
589 general, and our method in particular. We consider two optimization levels: the standard `-O0` (no
590 optimization at all), and the ad hoc customized `-O1` level (designed to limit CFG transformation
591 and maximize traceability). Since the LRA analysis is performed at the C level, the flow facts
592 discovered are the same whatever is the optimization level. A lack of improvement in the case of
593 optimized code is then necessarily du to an “imperfect” traceability.

594 The first group of columns recalls the results obtained with `-O0`; it only gives the best result,
595 obtained for one of the possible abstract domains (refer to Table 6 for details). The second group
596 gives information on the optimized code:

¹⁰Results obtained on an Intel(R) Xeon(R) CPU E5-2683 v3 @ 2.00GHz

- 597 ■ the initial WCET estimation given by OTAWA, together with the corresponding *speedup* factor
- 598 which indicates how “faster” is the optimized code compared to the non-optimized one;
- 599 ■ the best WCET estimation (together with the improvement percentage) obtained using PAGAI;
- 600 ■ the *traceability* ratio indicates how many counters introduced by our method are actually
- 601 associated to some basic block in the binary code. With a traceability of 100%, we expect
- 602 to observe an improvement percentage of the same order than the one obtained on the
- 603 non-optimized code. Note that the traceability with the non-optimized code is always 100%.

■ **Table 5** TacleBench functions Reference Labels

Ref	Directory	Function Names
ad.6	sequential/adpcm_dec	adpcm_dec_logsch
ad.7	sequential/adpcm_dec	adpcm_dec_logschl
ad.14	sequential/adpcm_dec	adpcm_dec_uppol2
ae.7	sequential/adpcm_enc	adpcm_enc_logsch
ae.8	sequential/adpcm_enc	adpcm_enc_logschl
ae.10	sequential/adpcm_enc	adpcm_enc_quantl
ae.16	sequential/adpcm_enc	adpcm_enc_uppol2
am.12	sequential/ammunition	ammunition_bit_string_set
am.17	sequential/ammunition	ammunition_divide_unsigned_integer
am.18	sequential/ammunition	ammunition_divide_unsigned_integer_without_overflow
am.47	sequential/ammunition	ammunition_multiply_integer
am.49	sequential/ammunition	ammunition_multiply_unsigned_integer
am.50	sequential/ammunition	ammunition_multiply_unsigned_integer_without_overflow
am.68	sequential/ammunition	ammunition_unsigned_integer_remainder
an.0	sequential/anagram	anagram_AddWords
an.2	sequential/anagram	anagram_BuildWord
an.8	sequential/anagram	anagram_init
an.14	sequential/anagram	anagram_ReadDict
an.15	sequential/anagram	anagram_Reset
an.16	sequential/anagram	anagram_return
bs.0	kernel/bsort	bsort_BubbleSort
bs.3	kernel/bsort	bsort_main
bs.4	kernel/bsort	bsort_return
bs.5	kernel/bsort	main
cr.2	crc	main
du.2	test/duff	duff_init
du.5	test/duff	main
ex.2	expint	main
gd.4	sequential/gsm_dec	gsm_dec_Coefficients_0_12
gd.5	sequential/gsm_dec	gsm_dec_Coefficients_13_26
gd.6	sequential/gsm_dec	gsm_dec_Coefficients_27_39
gd.11	sequential/gsm_dec	gsm_dec_Decoding_of_the_coded_Log_Area_Ratios
gd.16	sequential/gsm_dec	gsm_dec_Postprocessing
ge.11	sequential/gsm_encode	Gsm_Preprocess
ge.13	sequential/gsm_encode	Gsm_Short_Term_Analysis_Filter
gs.2	sequential/g723_enc	g723_enc_fmilt
gs.8	sequential/g723_enc	g723_enc_predictor_pole
gs.9	sequential/g723_enc	g723_enc_predictor_zero
gs.10	sequential/g723_enc	g723_enc_quan
gs.11	sequential/g723_enc	g723_enc_quantize
gs.16	sequential/g723_enc	g723_enc_update
hd.1	sequential/h264_dec	h264_dec_init
lc.0	lcdnum	main
li.3	app/lift	lift_controller
li.7	app/lift	lift_ctrl_set_vals
md.3	kernel/md5	md5_final
md.5	kernel/md5	md5_InitRandomStruct
md.13	kernel/md5	md5_R_RandomInit
md.14	kernel/md5	md5_R_RandomUpdate
md.15	kernel/md5	md5_transform
md.16	kernel/md5	md5_update
mp.9	sequential/mpeg2	mpeg2_frame_estimate
mp.11	sequential/mpeg2	mpeg2_fullsearch
sh.2	kernel/sha	sha_final
sm.0	sequential/statemate	main
sm.1	sequential/statemate	statemate_FH_DU
sm.2	sequential/statemate	statemate_generic_BLOCK_ERKENNUNG_CTRL
sm.3	sequential/statemate	statemate_generic_EINKLEMMSCHUTZ_CTRL
sm.4	sequential/statemate	statemate_generic_FH_TUERMODUL_CTRL
sm.8	sequential/statemate	statemate_main

■ **Table 6** How LRA can improve the estimated WCET of TacleBench

Ref	Initial WCET	Box			Octagons			Polyhedra		
		Δ	Imp ^t	Time	Δ	Imp ^t	Time	Δ	Imp ^t	Time
md.13	2648	0	0.0	<1s	1920	72.5	<1s	1920	72.5	<1s
an.15	173K	0	0.0	<1s	121K	69.9	1s	121K	69.9	<1s
gs.2	1105	0	0.0	<1s	738	66.7	9s	738	66.7	4s
hd.1	2092K	0	0.0	<1s	1371K	65.5	<1s	1371K	65.5	<1s
gs.8	2268	0	0.0	<1s	1476	65.0	1m	1476	65.0	1m
gs.9	6934	0	0.0	1s	738	10.6	4m	4428	63.8	29m
du.2	19K	0	0.0	<1s	12K	60.1	<1s	12K	60.1	<1s
du.5	22K	0	0.0	<1s	12K	53.8	<1s	12K	53.8	<1s
mp.9	52M	27M	51.1	56m	-	-	-	-	-	-
mp.11	10M	5M	51.1	2m	-	-	-	-	-	-
cr.2	227K	111K	48.7	<1s	111K	48.7	4s	111K	48.7	1s
lc.0	1540	0	0.0	<1s	595	38.6	<1s	595	38.6	<1s
md.15	8600	0	0.0	<1s	2304	26.7	<1s	2304	26.7	<1s
an.2	235K	0	0.0	<1s	58K	24.8	16s	58K	24.8	2s
an.0	466M	0	0.0	4s	4M	0.8	7m	115M	24.6	1m
md.5	13M	0	0.0	2m	-	-	-	3M	21.0	35m
ex.2	278K	1K	0.1	4s	28K	9.9	23s	53K	19.2	6s
sm.3	87	16	18.3	<1s	16	18.3	<1s	16	18.3	<1s
md.3	34K	0	0.0	18s	6K	17.2	6m	6K	18.2	31m
md.16	13K	0	0.0	7s	2K	17.7	19s	2K	17.7	10s
sm.0	268K	46K	17.1	16s	-	-	-	-	-	-
gs.10	942	0	0.0	<1s	126	13.3	1s	126	13.3	<1s
am.47	3649	0	0.0	13m	76	2.0	14m	485	13.2	14m
gs.11	2020	0	0.0	1s	252	12.4	17s	252	12.4	4s
md.14	51K	0	0.0	2m	5K	10.4	21m	-	-	-
sm.4	595	62	10.4	3s	62	10.4	2m	62	10.4	72m
li.7	1088	0	0.0	<1s	102	9.3	2s	102	9.3	<1s
gs.16	3760	0	0.0	7s	254	6.7	63m	-	-	-
ad.6	66	0	0.0	<1s	4	6.0	<1s	4	6.0	<1s
ae.7	66	0	0.0	<1s	4	6.0	<1s	4	6.0	<1s
ad.7	67	0	0.0	<1s	4	5.9	<1s	4	5.9	<1s
ae.8	67	0	0.0	<1s	4	5.9	<1s	4	5.9	<1s
sm.1	266K	14K	5.1	21s	-	-	-	-	-	-
sm.8	266K	14K	5.1	26s	-	-	-	-	-	-
an.16	530	0	0.0	<1s	25	4.7	<1s	25	4.7	<1s
am.50	2251	0	0.0	15m	76	3.3	15m	95	4.2	15m
am.49	2281	0	0.0	9m	76	3.3	14m	95	4.1	12m
sm.2	248	10	4.0	<1s	10	4.0	1s	10	4.0	1s
bs.4	5580	0	0.0	<1s	196	3.5	<1s	196	3.5	<1s
am.12	468	0	0.0	12m	16	3.4	12m	16	3.4	12m
ad.14	120	0	0.0	<1s	4	3.3	<1s	4	3.3	<1s
ae.16	120	0	0.0	<1s	4	3.3	<1s	4	3.3	<1s
li.3	3405	0	0.0	11s	102	2.9	47m	-	-	-
ae.10	1473	0	0.0	<1s	33	2.2	<1s	33	2.2	<1s
ge.11	47K	0.16K	0.3	1s	1K	2.0	1m	1K	2.0	16s
am.68	12K	0	0.0	10m	0.15K	1.3	73m	-	-	-
gd.16	23K	0.32K	1.3	<1s	0.32K	1.3	2s	0.32K	1.3	<1s
gd.4	1333	16	1.2	1s	16	1.2	6s	16	1.2	1s
gd.6	1333	16	1.2	<1s	16	1.2	3s	16	1.2	<1s
sh.2	25K	0.31K	1.2	6s	0.31K	1.2	3m	0.31K	1.2	1m
an.8	3079K	0	0.0	<1s	0	0.0	7s	34K	1.1	1s
an.14	3079K	0	0.0	<1s	0	0.0	5s	34K	1.1	1s
gd.11	1420	16	1.1	1s	16	1.1	35m	-	-	-
ge.13	727K	8K	1.0	1m	-	-	-	-	-	-
bs.0	1045K	0	0.0	<1s	0	0.0	2s	10K	0.9	<1s
bs.3	1045K	0	0.0	<1s	0	0.0	1s	10K	0.9	<1s
bs.5	1053K	0	0.0	<1s	0.20K	0.0	7s	10K	0.9	1s
gd.5	837	8	0.9	<1s	8	0.9	1s	8	0.9	<1s
am.17	8930	0	0.0	9m	76	0.8	20m	-	-	-
am.18	8901	0	0.0	11m	76	0.8	20m	-	-	-

■ **Table 7** Loop bounds discovery, using PAGAI without the help of oRange nor the user-given bounds. This experiment is performed for the 54 cases from Table 6 where PAGAI terminates with either octagons or polyhedra ; the box domain is unable to find loop bounds and is not considered here. Within this test set, 10 programs contain no loop and are thus trivially bounded (adVI, adVII, adXIV, aeVII, aeVIII, aeXVI, gdXI, smII, smIII, smIV). For the remaining programs, first column gives the depth of nested loops and column two indicates if PAGAI gives a bounded (i.e., finite) WCET estimation. The WCET value is not given: it corresponds to the best PAGAI estimation in Table 6. The “paradoxal” result for ex.2 (loop depth 2 and bounded) is due to the fact that PAGAI “bounds” the inner-loop to 0 (i.e., the loop appears in a infeasible branch).

	loop depth	PAGAI
ae.10	1	bounds
bs.4	1	bounds
gd.4	1	bounds
gd.5	1	bounds
gd.6	1	bounds
gd.16	1	bounds
ge.11	1	bounds
gs.2	1	bounds
gs.8	1	bounds
gs.16	1	bounds
lc.0	1	bounds
li.7	1	bounds
md.15	1	bounds
du.2	1	bounds
du.5	1	bounds
hd.1	1	bounds
md.13	1	bounds
an.15	1	bounds
an.16	1	bounds
am.12	1	T
an.2	1	T
gs.10	1	T
gs.11	1	T
li.3	1	T
sh.2	1	T

	loop depth	PAGAI
ex.2	2	bounds
am.47	2	T
am.49	2	T
am.50	2	T
an.8	2	T
an.14	2	T
bs.0	2	T
bs.3	2	T
bs.3	2	T
cr.2	2	T
gs.9	2	T
md.3	2	T
md.14	2	T
md.16	2	T
am.17	3	T
am.18	3	T
am.68	3	T
an.0	3	T
md.5	4	T

■ **Table 8** Observing the impact of compilation levels on LRA

Ref	OO			Opt. speedup	CO			Traceability	
	Initial WCET	Best WCET	Best Imp ^t		Initial WCET	Best WCET	Best Imp ^t		
md.13	2648	728	72.5	3.3x	791	215	72.8	100%	of 2
an.15	173K	52K	69.9	3.0x	58K	17K	69.9	100%	of 6
gs.2	1105	367	66.7	2.3x	479	171	64.3	100%	of 14
hd.1	2092K	721K	65.5	2.1x	1019K	350K	65.6	100%	of 4
gs.8	2268	792	65.0	2.4x	963	347	63.9	100%	of 28
gs.9	6934	2506	63.8	2.4x	2910	1062	63.5	100%	of 28
du.2	19K	8K	60.1	2.3x	8K	3K	64.0	100%	of 2
du.5	22K	10K	53.8	2.4x	9K	4K	59.3	100%	of 3
mp.9	52M	25M	51.1	4.6x	11M	11M	0.0	79%	of 890
mp.11	10M	5M	51.1	4.6x	2M	2M	0.0	79%	of 178
cr.2	227K	116K	48.7	2.3x	97K	50K	48.7	41%	of 24
lc.0	1540	945	38.6	2.4x	641	421	34.3	100%	of 4
md.15	8600	6296	26.7	2.8x	3064	2200	28.1	100%	of 2
an.2	235K	176K	24.8	2.9x	80K	59K	25.9	100%	of 21
an.0	466M	351M	24.6	3.0x	157M	116M	25.9	100%	of 44
md.5	13M	10M	21.0	3.4x	4M	3M	15.0	80%	of 40
ex.2	278K	224K	19.2	1.3x	218K	218K	0.0	46%	of 13
sm.3	87	71	18.3	1.1x	78	59	24.3	100%	of 5
md.3	34K	28K	18.2	2.7x	13K	10K	17.1	100%	of 23
md.16	13K	11K	17.7	2.6x	5K	4K	17.2	100%	of 10
sm.0	268K	222K	17.1	1.1x	237K	193K	18.8	100%	of 108
gs.10	942	816	13.3	2.5x	381	325	14.6	100%	of 4
am.47	3649	3164	13.2	2.6x	1417	1282	9.5	100%	of 26
gs.11	2020	1768	12.4	2.4x	830	718	13.4	100%	of 11
md.14	51K	46K	10.4	2.7x	19K	17K	10.8	97%	of 37
sm.4	595	533	10.4	1.1x	547	493	9.8	100%	of 42
li.7	1088	986	9.3	2.1x	516	468	9.3	100%	of 9
gs.16	3760	3506	6.7	2.3x	1653	1539	6.8	100%	of 69
ad.6	66	62	6.0	3.0x	22	20	9.0	100%	of 3
ae.7	66	62	6.0	3.0x	22	20	9.0	100%	of 3
ad.7	67	63	5.9	2.9x	23	21	8.6	100%	of 3
ae.8	67	63	5.9	2.9x	23	21	8.6	100%	of 3
sm.1	266K	252K	5.1	1.1x	237K	224K	5.1	100%	of 97
sm.8	266K	252K	5.1	1.1x	237K	224K	5.1	100%	of 96
an.16	530	505	4.7	1.7x	316	299	5.3	100%	of 3
am.50	2251	2156	4.2	2.6x	860	823	4.3	100%	of 8
am.49	2281	2186	4.1	2.6x	865	839	3.0	100%	of 9
sm.2	248	238	4.0	1.1x	230	220	4.3	100%	of 11
bs.4	5580	5384	3.5	1.9x	2883	2687	6.7	100%	of 5
am.12	468	452	3.4	2.6x	181	177	2.2	100%	of 13
ad.14	120	116	3.3	2.6x	46	44	4.3	100%	of 8
ae.16	120	116	3.3	2.6x	46	44	4.3	100%	of 8
li.3	3405	3303	2.9	1.6x	2093	2045	2.2	100%	of 57
ae.10	1473	1440	2.2	2.1x	706	690	2.2	100%	of 6
ge.11	47K	46K	2.0	2.5x	19K	18K	5.0	100%	of 24
am.68	12K	12K	1.3	1.2x	9K	9K	0.5	100%	of 52
gd.16	23K	23K	1.3	2.3x	10K	10K	3.1	100%	of 12
gd.4	1333	1317	1.2	2.4x	553	537	2.8	100%	of 12
gd.6	1333	1317	1.2	2.4x	553	537	2.8	100%	of 12
sh.2	25K	24K	1.2	2.2x	11K	11K	0.9	76%	of 39
an.8	3079K	3045K	1.1	2.3x	1333K	1310K	1.7	100%	of 14
an.14	3079K	3045K	1.1	2.3x	1333K	1310K	1.7	100%	of 15
gd.11	1420	1404	1.1	2.4x	601	585	2.6	100%	of 121
ge.13	727K	719K	1.0	1.9x	377K	369K	2.0	100%	of 289
bs.0	1045K	1035K	0.9	2.7x	389K	385K	0.9	100%	of 6
bs.3	1045K	1035K	0.9	2.7x	389K	385K	0.9	100%	of 5
bs.5	1053K	1043K	0.9	2.7x	393K	389K	1.0	100%	of 10
gd.5	837	829	0.9	2.5x	337	329	2.3	100%	of 7
am.17	8930	8854	0.8	1.1x	8431	8405	0.3	100%	of 36
am.18	8901	8825	0.8	1.1x	8426	8400	0.3	100%	of 36