

Available online at www.sciencedirect.com

The Journal of Systems and Software xxx (2004) xxx–xxx

www.elsevier.com/locate/jss

2 *ConSUS*: a light-weight program conditioner

3 Sebastian Danicic^{a,*}, Mohammed Daoudi^a, Chris Fox^b, Mark Harman^c,
4 Rob M. Hierons^c, John R. Howroyd^a, Lahcen Ourabya^a, Martin Ward^d

5 ^a Department of Mathematics and Computer Science, Goldsmiths College, University of London, New Cross, London SE14 6NW, United Kingdom

6 ^b Department of Computer Science, University of Essex, Wivenhoe Park, Colchester, CO4 3SQ, United Kingdom

7 ^c Department of Information Systems and Computing, Brunel University, Uxbridge, Middlesex, UB8 3PH, United Kingdom

8 ^d Software Technology Research Laboratory, De Montfort University, The Gateway, Leicester LE1 9BH, United Kingdom

Received 1 April 2003; received in revised form 16 July 2003; accepted 2 March 2004

11 Abstract

12 Program conditioning consists of identifying and removing a set of statements which cannot be executed when a condition of
13 interest holds at some point in a program. It has been applied to problems in maintenance, testing, re-use and re-engineering.
14 All current approaches to program conditioning rely upon both symbolic execution and reasoning about symbolic predicates.
15 The reasoning can be performed by a ‘heavy duty’ theorem prover but this may impose unrealistic performance constraints.

16 This paper reports on a lightweight approach to theorem proving using the FermaT simplify decision procedure. This is used as a
17 component to *ConSUS*, a program conditioning system for the Wide Spectrum Language WSL. The paper describes the symbolic
18 execution algorithm used by *ConSUS*, which prunes as it conditions.

19 The paper also provides empirical evidence that conditioning produces a significant reduction in program size and, although
20 exponential in the worst case, the conditioning system has low degree polynomial behaviour in many cases, thereby making it scal-
21 able to unit level applications of program conditioning.

22 © 2004 Published by Elsevier Inc.

23 *Keywords*: Program conditioning; Slicing; Program transformation; Decision procedures

25 1. Introduction

26 Program conditioning¹ (Canfora et al., 1998; Har-
27 man et al., 2001), like program slicing (Weiser, 1984),
28 is a form of source code manipulation that allows a soft-
29 ware engineer to extract an executable sub-program
30 based upon a criterion of interest. The original formula-
31 tion of slicing (Weiser, 1984) was static. That is, the slic-
32 ing criterion contained no information about the input
33 to the program. A static end-slice of program P with re-
34 spect a set of variables V , is a program P' that ‘behaves

the same’ as P with respect to all the variables in V . Fur- 35
thermore, P' is obtained from P by statement deletion. 36
The way in which slicing produces an executable sub- 37
program, based upon some criterion of interest, gives 38
rise to many applications. For example, slicing has been 39
applied to, among others, debugging (Lyle and Weiser, 40
1987; Weiser, 1982), testing (Binkley, 1998; Harman 41
and Danicic, 1995; Hierons et al., 1999; Hierons et al., 42
2002), program comprehension (Binkley et al., 2000; 43
Fox et al., 2001), program decomposition (Gallagher 44
and Lyle, 1991) and integration (Binkley et al., 1995; 45
Horwitz et al., 1989), software metrics (Bieman and 46
Ott, 1994; Ott and Thuss, 1989) and re-engineering 47
and reverse engineering (Canfora et al., 1994a,b; Cimi- 48
tile et al., 1995a,b). 49

* Corresponding author..

E-mail address: s.danicic@gold.ac.uk (S. Danicic).

¹ Henceforth, the phrase ‘program conditioning’ will be referred to simply as ‘conditioning’.

Static slicing has now reached a mature stage of development, in which tools such as the Wisconsin Program Slicing System, marketed through Grammatech (Grammatech Inc., 2002) can efficiently slice real-world C programs of the order of hundreds of thousands of lines of code in reasonable time (Binkley and Harman, to appear-a). Conditioning is central in a variation of slicing called conditioned slicing.

Conditioned slicing forms a theoretical bridge between the two extremes of static and dynamic slicing. It augments the traditional slicing criterion with a condition which captures a set of initial program states of interest. This additional condition can be used to simplify the program before applying a traditional static slicing algorithm. Such pre-simplification is called conditioning, and it is achieved by eliminating statements which do not contribute to the computation of the variables of interest when the program is executed in an initial state which satisfies the condition.

Conditioning is defined independently of slicing: the result of conditioning P will be a program, Q , that behaves the same as P whenever the inputs satisfy a path condition. A path condition is simply a boolean expression involving some or all of the program variables. In the approach described in this paper conditions of interest are expressed as Assert statements, i.e. boolean expressions, which may be arbitrarily placed throughout the program. In this case, the resulting program, Q , must agree with original, P , for all inputs where P satisfies these intermediate Assert statements. Even where no Assert statements are added, the system will attempt to remove infeasible paths (a useful step in itself).

The paper focuses upon the conditioning step in producing a conditioned slice. This is because the slicing step is standard static slicing, for which a WSL static slicer is used, which implements an extension (Ouarbya et al., 2002) of Hausler's data-flow based approach to static slicing (Hausler, 1989). The static slicing phase is not discussed further. It is the conditioning phase that makes conditioned slices interesting, smaller than their static counter-parts (in general) and, crucially, which increases the difficulty involved in slice construction (because both symbolic execution and some form of theorem proving is required).

As an example of the way in which conditioning identifies sub-programs, consider the Taxation program in Fig. 1. The figure contains a fragment² of a program which encodes the UK tax regulations in the tax year April 1998–April 1999. Each person has a personal allowance which is an amount of un-taxed income. The size of this personal allowance depends upon the status of the person, which is encoded in the boolean

variables `blind`, `married` and `widowed`, and the integer variable `age`. For example, given the condition

$$\text{age} \geq 65 \text{ AND } \text{age} < 75 \text{ AND } \text{income} = 36000 \text{ AND } \text{blind} = 0 \text{ AND } \text{married} = 1$$

conditioning the program identifies the statements which appear boxed in the figure. This is useful because it allows the software engineer to isolate a sub-computation concerned with the initial condition of interest. The sub-program extracted can be compiled and executed as a separate code unit. It will be guaranteed to mimic the behaviour of the original if the initial condition is met.

In the worst case, there is no doubt that the time to perform 'best possible' program conditioning will be exponential in the size of program being conditioned. This fact is inherent to the problem of conditioning. It can be seen that by considering a sequence of simple IF-THEN-ELSE statements. The number of paths is clearly 2^n . If no pruning is possible, then the number of calls to the theorem prover will therefore be $O(2^n)$. The same, however, is also true of many aspects of theorem proving. For example, the Boolean Satisfiability Problem, which the CHAFF (Moskewicz et al., 2001) solver at the heart of the CVC theorem prover (Stump et al., 2002) implements a solution to, is a well known NP-Complete problem. As in the case of theorem provers, this worst case scenario does not imply that implementations of conditioners are infeasible. The reasons for this are twofold:

- (1) In conditioning, *any* resulting reduction in program size represents progress. Even if the best possible results require exponential time, it is possible that significant reductions will be performed more quickly.
- (2) In many cases conditioning may be performed in low order polynomial or even quadratic time as suggested by the empirical study in this paper. In such cases conditioning may be applicable to unit level applications at least.

The language of implementation and the language which is sliced by the approach reported here is WSL, the Wide Spectrum Language, introduced for reverse engineering () through program transformation (Ward, 1989, 1994, 1999). WSL uses an Algol-like syntax, but has additional facilities to make it wide-spectrum and to allow transformations to be expressed within WSL itself.

The contributions of this paper are:

- To define a new more efficient algorithm for program conditioning.
- To report on empirical studies which demonstrate that:

² This is WSL version of the C program previously used in Danicic et al. (2000).

```

IF (age>=75) THEN personal := 5980
ELSE IF (age>=65)
    THEN personal := 5720
    ELSE personal := 4335
    FI
FI;
IF (age>=65 AND income >16800)
THEN IF (4335 > personal-((income-16800) / 2))
    THEN personal := 4335
    ELSE personal := personal-((income-16800) / 2)
    FI
FI;
IF (blind =1) THEN personal := personal + 1380 FI;
IF (married=1 AND age >=75)
THEN pc10 := 6692
ELSE IF (married=1 AND age >= 65)
    THEN pc10 := 6625
    ELSE IF (married=1 OR widow=1)
        THEN pc10 := 3470
        ELSE pc10 := 1500
        FI
    FI
FI;
IF (married=1 AND age >= 65 AND income > 16800)
THEN
    IF (3470 > pc10-(income-16800) / 2)
    THEN pc10 :=3470
    ELSE pc10 := pc10-((income-16800) / 2)
    FI
FI;
IF (income - personal <= 0)
THEN tax := 0
ELSE income := income - personal;
FI;
IF (income <= pc10)
THEN tax := income * rate10
ELSE tax := pc10 * rate10;
    income := income - pc10;
FI;
IF (income <= 28000)
THEN tax := tax + income * rate23
ELSE tax := tax + 28000 *rate23;
    income := income - 28000;
    tax := tax + income * rate40
FI

```

Conditioned program: boxed lines of code

Condition: age>=65 AND age<75 AND income=36000
AND blind=0 AND married=1

Fig. 1. A fragment of the taxation calculation program in WSL.

- 154 (a) On small 'real programs' *ConSUS* produces a
considerable reduction in program size when
used with and without a program slicer.
- 157 (b) The *ConSUS* algorithm when used in conjunc-
tion with WSL's *FermaT Simplify* has the
161 potential for 'scaling up' for use on large systems.

The rest of this paper is organised as follows: related 162
work is discussed in Section 2. Section 3 introduces and
defines conditioning. In Section 4, the new conditioning
algorithm is defined in detail. In Section 5, our empirical
studies are presented. Conclusions and directions for fu-
ture work are presented in Section 6.

2. Related work

Program slicing originated with Mark Weiser's doctoral thesis (Weiser, 1979). Weiser's formulation of slicing was a static one, captured by the slicing criterion, which was a set of variables and a program point (Weiser, 1984). This static paradigm for slicing, was later augmented by a dynamic paradigm (Korel and Laski, 1988). Dynamic slicing represented the first move away from the static paradigm, indicating that there may be other ways to formulate a slicing criterion.

The move from static to other paradigms was further developed with the introduction of a form of slicing called quasi-static slicing which bridged the gap between static and dynamic slicing (Venkatesh, 1991). A quasi-static slice is constructed with respect to a prefix of input, rather than a full input sequence. Quasi-static slicing was motivated by work on partial evaluation (Ershov, 1978; Futamura, 1971). The concept of quasi-static slicing is subsumed by conditioned slicing, which is the most general form of slicing, hitherto explored in the slicing literature.

Conditioned slicing, the subject of this paper, was introduced by Canfora et al. (Canfora et al., 1994a) in 1994. Field et al. (1995) introduced a similar technique called constrained slicing, which also uses conditions to specialize programs. In constrained slicing, parts of the program which cannot contribute to the values of variables of interest are identified as holes. Because it contains holes, a constrained slice is not necessarily an executable subprogram, as a conditioned slice is.

Earlier, ideas and techniques very similar to program conditioning were introduced in Coen-Porisini et al. (1991). This work uses symbolic execution and theorem proving to specialize Ada programs. In their approach, generalised software components are specialised by restricting their domain of inputs thereby increasing the efficiency of components for each particular instance of their use.

To implement conditioned slicing it is necessary to use some form of theorem proving technology. In early work on conditioned slicing, the theorem proving was performed by hand (Canfora et al., 1994a, 1998; De Lucia et al., 1996), or by term-rewriting (in the case of constrained slicing, Field et al., 1995). More recent work explored the use of Isabelle (Danicic et al., 2000) and SVC (Danicic et al., accepted). This work using Isabelle and SVC can be thought of as a proof-of-concept implementation, showing how the theorem prover could be interchanged as a component of the overall conditioning approach. The present work explores the use of the FermaT simplify transformation, a decision procedure which is used to implement a form of highly light-weight theorem proving.

The paper also introduces a new approach to symbolic execution which follows a different form to previ-

ous symbolic executors (Danicic et al., accepted; King, 1976; Coen-Porisini and De Paoli, 1990; Coen-Porisini et al., 1991) in its handling of loops and which is more closely integrated with the theorem prover than previous approaches (Canfora et al., 1998; Danicic et al., 2000; Danicic et al., accepted),³ thereby exploiting the possibility of path-pruning to speed up the conditioning process.

Other authors have considered the way in which theorem provers may be used in slicing (Krinke and Snelting, 1998) and the way in which forms of symbolic execution can assist related analyses such as constrained test data generation (DeMillo and Offutt, 1993; King and Offutt, 1991; Offutt, 1990).

In the VALSoft project (Krinke and Snelting, 1998), Krinke and Snelting show how the computation of symbolic values can be used to refine data-flow relations between array assignments and uses. Specifically, they consider the situation where an assignment to an array like

```
A[i] := 4;
```

may not filter through to a reference to the same array 245

```
IF A[j] == K THEN ... 247
```

because it can be statically determined that i and j are 248
guaranteed to be non-equal at the `IF` statement. In order 249
to determine this statically, Krinke and Snelting 250
use a simple form of symbolic execution and then a theorem 251
prover to determine (statically) the properties of 252
expression indices at array expressions. This work differs 253
from the work reported here because the application of 254
theorem prover and symbolic executor is used to refine 255
the *static* data dependence relation used to construct a 256
static slice. 257

In DeMillo and Offutt (1993), King and Offutt (1991) 258
and Offutt (1990), constraint solving techniques are used 259
to support automated test data generation. This ap- 260
proach requires the backward propagation of con- 261
straints from a point of interest, p , within the 262
program. This is a form of backward symbolic execu- 263
tion. The constraints are (ideally) propagated back to 264
the starting state of the program, where they become 265
constraints on the input space. The solution of these in- 266
put constraints is a valid test data vector to execute 267
point p in the desired way. Constraint solving can be 268
achieved using a theorem prover or a simple decision 269
procedure, offering similar speed/precision trade-offs to 270
those considered in the present paper. 271

There are several surveys of slicing in the literature 272
which cover dynamic slicing and its applications (Korel 273
and Rilling, 1998), slicing techniques and algorithms 274
(Tip, 1995), forms of slicing and their applications 275

³ It has some similarities to the approach taken by Coen-Porisini et al. in Coen-Porisini et al. (1991). See Section 4.

276 (Binkley and Gallagher, 1996; De Lucia, 2001; Harman
 277 and Hierons, 2001) and empirical results on the applica-
 278 tion of slicing (Binkley and Harman, to appear-b).

279 **3. Conditioning**

280 Conditioning is the act of simplifying a program
 281 assuming that the states of the program at certain cho-
 282 sen points in its execution satisfy certain properties.
 283 These properties of interest are expressed by adding As-
 284 sert statements to the program being conditioned. Con-
 285 sider, for example the program in Fig. 2. Here, program
 286 simplification is being attempted assuming that it is exe-
 287 cuted in an initial state where $x > y$. In such states, the
 288 *true* path of the IF-THEN-ELSE statement will always
 289 be taken and thus the program can be simplified to
 290 $\{x > y\}; a := 1$. Notice that the Assert statement is also
 291 included in the resulting conditioned program. This is
 292 because an Assert statement is a valid WSL statement
 293 that aborts if its condition is *false*. Observe that the origi-
 294 nal program's behaviour and that of the conditioned
 295 program are identical.

296 A software engineer may require a program to be
 297 conditioned with respect to intermediate states as well
 298 as with respect to initial states. The example in Fig. 3 ex-
 299 presses the fact that the program is to be simplified
 300 assuming that all its inputs are positive. The conditioner
 301 should be able to replace the final IF-THEN-ELSE state-
 302 ment by $\text{PRINT}(\backslash\text{POSITIVE})$.

303 A conditioner is a program which tries to remove
 304 code that is unreachable given the assertions. Therefore,
 305 a conditioner will try to remove unreachable code even if
 306 the program contains no Assert statements (see Fig. 4
 307 for an example of this).

308 Conditioners are required to reason about the valid-
 309 ity of paths under certain conditions. In order to per-
 310 form such reasoning, it appears sensible to utilise
 311 existing automated theorem provers rather than to de-

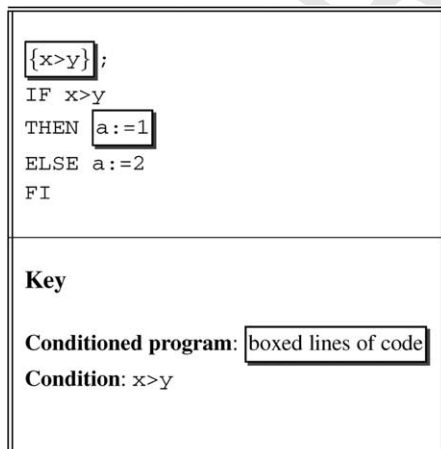


Fig. 2. Conditioning a simple program.

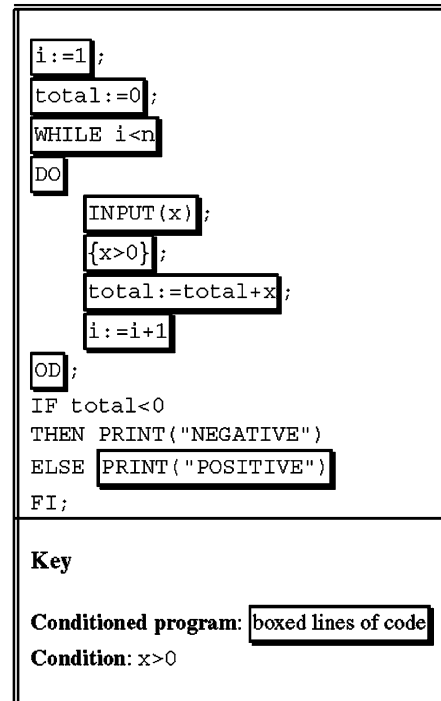


Fig. 3. Conditioning with Intermediate Asserts.

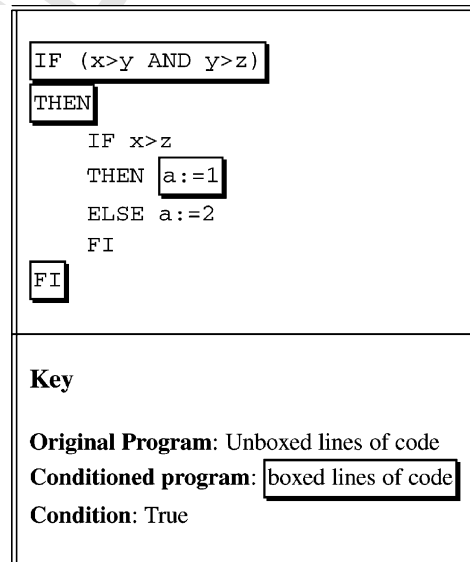


Fig. 4. Conditioning without assert.

336 velop new ones. Consider the program in Fig. 5. Here, 336
 337 conditioning of a simple IF statement assuming that 337
 338 the initial state has the property that $x > y \wedge y > z$ is 338
 339 being attempted. This is achieved by adding the corre- 339
 340 sponding Assert statement at the beginning of the pro- 340
 341 gram. The simplification achieved depends upon the 341
 342 conditioner's ability to infer that $x > y \wedge y > z \Rightarrow x > z$. 342
 343 If the conditioner knows that the operator, $>$, is transi- 343
 344 tive, then it will be able to infer that the second of these 344

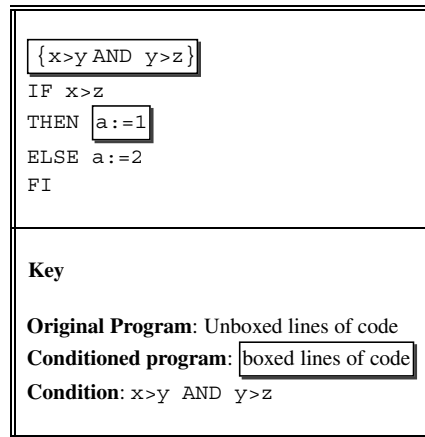


Fig. 5. Conditioning a simple program.

345 conditions is a contradiction and therefore that the
 346 ELSE branch of the IF is infeasible. Only the Assert
 347 statement, $\{x > y$ AND $y > z\}$, and assignment $a := 1$ are
 348 required; the rest of the code can be removed. The sim-
 349 plifying power of the conditioner depends on two things:
 350 (1) The precision of the symbolic executor which han-
 351 dles propagation of state and path information.
 352 (2) The precision of the underlying theorem prover
 353 which determines the truth of propositions about
 354 states and paths.

355 By using an approximation to a program's semantics
 356 using a form of symbolic execution, and by being willing
 357 to accept approximate results from the theorem proving
 358 itself, conditioning allows us to adopt reasoning that
 359 does not require the full force of inductive proofs. The
 360 theorem proving used in programming conditioning is
 361 lightweight when compared to the theorem proving re-
 362 quired for a complete formal analysis of a program.
 363 The problem can be further constrained to cases where
 364 the theorem proving can be implemented by completed
 365 decision procedures. There are limitations to the kinds
 366 of expressions for which complete decision procedures
 367 exist, one typical limitation is a restriction to reasoning
 368 with sets of so-called 'linear' inequalities.

370 4. The ConSUS conditioning algorithm

371 4.1. An overview of the approach

372 When implementing an interpreter, a program is eval-
 373 uated in a state, which maps variables to their values
 374 (Stoy, 1985). In symbolic execution (Coen-Porisini and
 375 De Paoli, 1990; Coward, 1988; Coward, 1991; Girgis,
 376 1992), the state, called a *symbolic store*⁴, maps varia-

bles, not to *values*, but to *symbolic expressions*. The
 expressions are the objects that can occur on the right
 hand sides of assignment statements (in this paper, it
 is assumed that these are simply arithmetic expressions).

When a program is symbolically evaluated in an ini-
 tial symbolic store it gives rise to a set of possible final
 symbolic stores. The reason that a symbolic evaluator
 returns a *set* of final stores is that our program may have
 more than one path, each of which may define a different
 final symbolic store. Unlike the case of an interpreter,
 the initial symbolic store does not give rise to a unique
 path through the program. A *symbolic evaluator* can,
 thus, be thought of as a mapping, which given a pro-
 gram and a symbolic store, returns a set of symbolic
 stores.

In order to implement a conditioner, a richer state
 space than that used in a symbolic evaluator is required.
 For each final symbolic store it is necessary also to re-
 cord what properties must have been true of the initial
 symbolic store in order for the program to take the path
 that resulted in this final symbolic store. This is called a
path condition and is simply a boolean expression involv-
 ing constants and variables of the program.

A *conditioned state*, Σ , is represented by a set of path
 condition, symbolic store pairs. A pair (b, σ) being an
 element of a conditioned state implies that the symbolic
 store σ can be reached if path condition b is true ini-
 tially. If a conditioned state contained the pair $(false, \sigma)$
 this is equivalent to stating that the symbolic store σ is
 unreachable.

ConSUS can be thought of as a function which takes
 a program and an initial conditioned state and returns a
 (simplified) program and a final conditioned state.⁵ In
 practice, a conditioner will normally be applied to pro-
 grams starting in the *natural conditioned state*. In the
 natural conditioned state, the corresponding symbolic
 store, *id* maps all variables to their names, representing
 the fact that no assignments have yet taken place. The
 corresponding path condition in the natural state is *true*,
 representing the fact that no paths have yet been taken.

4.1.1. Statement removal

The program simplification produced by *ConSUS*
 arises from the fact that a statement from a program
 can be removed if all paths, starting from the initial con-
 ditioned state of interest, leading to the statement are
 infeasible. The path condition corresponding to a sym-
 bolic store is a condition which must be satisfied by
 the initial store in order for the program to take the path
 that arrives at the corresponding symbolic store. If,
 therefore, the final path condition, is equivalent to false

⁴ Usually called the *symbolic state*.

⁵ In (Coen-Porisini et al., 1991), similar functions *exec* and *simpl* are defined. Fundamentally different, however, is that *exec* and *simpl* return a single path condition, symbolic state pair, not a set of such pairs as in our case.

427 (a contradiction) this means that the store is not reach- 479
 428 able. If, on the other hand, the final path condition is 480
 429 equivalent to true (a tautology) then all paths starting 481
 430 from the given initial store will lead to the corresponding 482
 431 symbolic store. The power of a conditioner, in essence, 483
 432 depends on the ability to prove that the path conditions 484
 433 encountered are tautologies or contradictions. This is 485
 434 why a conditioner needs to work in conjunction with a 486
 435 theorem prover. Of course, to do this perfectly is not a 487
 436 computable problem and therefore, in general, infeasible 488
 437 paths will be unnecessarily considered. 489

438 Consider again, the program in Fig. 5. This program 490
 439 potentially has two possible final symbolic stores: 491

[$a \rightarrow 1$]

441 [$a \rightarrow 2$]

442 The corresponding path conditions are

$x > y \wedge y > z \wedge x > z$

444 $x > y \wedge y > z \wedge \neg(x > z)$.

445 Combining these two gives the conditioned state with 492
 446 two elements: 493

($x > y \wedge y > z \wedge x > z$, [$a \rightarrow 1$])

448 ($x > y \wedge y > z \wedge \neg(x > z)$, [$a \rightarrow 2$]).

449 A sufficiently powerful theorem prover will be able to in- 494
 450 fer that the second of these path conditions is a 495
 451 contradiction. 496

452 Often programs containing no Assert statements will 497
 453 be conditioned. This corresponds to removing *dead* 498
 454 code. Consider the program in Fig. 4. The programs 499
 455 in Figs. 5 and 4 do not quite have the same semantics. 500
 456 The first will abort in initial stores not satisfying the ini- 501
 457 tial path condition, while the second will do nothing but 502
 458 terminate successfully in these stores. The dead code 503
 459 $a:=2$ is removed by the conditioner in both cases. 504

460 As will be shown later, *ConSUS* is efficient in the 505
 461 sense that it attempts to prune paths ‘on the fly’ as it 506
 462 symbolically executes. This is clearly an improvement 507
 463 over other systems like *ConSIT* (Danicic et al., 2000) 508
 464 which generates all paths and then prunes once at the 509
 465 end. The way this is achieved, is that on encountering 510
 466 a guard, *ConSUS* interacts with its theorem proving 511
 467 mechanism to check whether the negation of the sym- 512
 468 bolic value of the guard is implied by the corresponding 513
 469 path condition in all values of the current conditioned 514
 470 state. If this is the case, then the corresponding body 515
 471 is unreachable and so can be removed without being 516
 472 processed.

473 Programs containing loops may have infinitely many 517
 474 paths. These cannot all be considered and therefore a 518
 475 conservative and safe approach has to be adopted when 519
 476 conditioning loops. For each WHILE loop, it is essential 520
 477 that in any implementation, only a finite number of dis- 521
 478 tinct symbolic stores are generated. A *meta symbolic*

store is required in order to represent the infinite set of 479
 symbolic stores that are not distinguished between. This 480
meta symbolic store must be safe in the sense that it must 481
 not add any untrue information about these symbolic 482
 stores. The simplest possible approach is simply to 483
 ‘throw away’ any information about variables which 484
 are affected by the body of a loop. This idea is very sim- 485
 ilar to state folding introduced in (Coen-Porisini et al., 486
 1991). Their program specializer, returns a single sym- 487
 bolic store, path condition pair, and so it is necessary 488
 to throw away values corresponding to variables as- 489
 signed different values on each branch of an IF THEN 490
 ELSE statement. 491

Using this approach, a WHILE loop will map each 492
 symbolic store, σ , to a set consisting of two symbolic 493
 stores. One of the stores will be σ itself, (representing 494
 the fact that the guard of the loop may be initially 495
false) and the other store (representing the fact that the 496
 loop was executed at least once) will be represented by a 497
 store, σ' , which agrees with σ on all variables not af- 498
 fected by the body of the loop. In σ' , all variables that 499
 are affected by the body of the loop are *skolemised*, rep- 500
 resenting that fact that we no longer have any informa- 501
 tion about their value. By skolemising a variable, all 502
 previous information that we had about it is being 503
 thrown away. As a result of skolemising a symbolic 504
 store, wrong information will never be generated, just 505
 less precise information. 506

The approach taken by *ConSUS* (based on the ap- 507
 proach of *ConSIT*) is less crude, however. In this case, 508
 as a result of symbolically evaluating a WHILE loop, 509
 there arises the set consisting of σ as before, together 510
 with the set of stores which are the result of symbolically 511
 executing the body of the loop in the skolemised store 512
 σ' . To see how these two approaches differ, consider 513
 the example given in Fig. 6. Using the naïve approach, 514
 the two symbolic stores resulting from the WHILE loop 515
 are [$x \rightarrow y + 1$] and [$x \rightarrow x_0$]. The first of these repre- 516

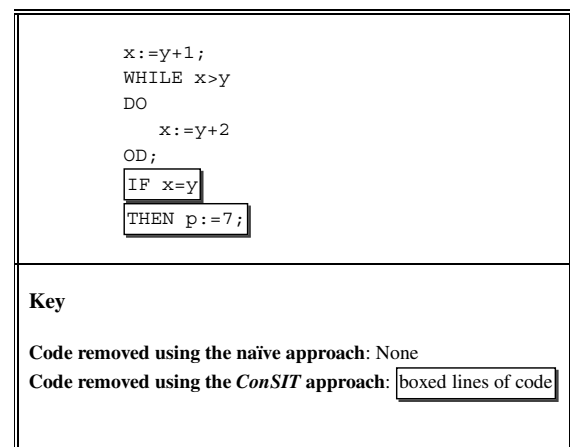


Fig. 6. Conditioning a WHILE loop using two approaches.

517 sends not executing the loop at all and the second repre-
 518 sents the fact that the loop body has been executed at
 519 least once. The variable x has been skolemised to x_0 ,
 520 representing the fact that its value is no longer known.
 521 Evaluating the guard $x = y$ of the IF–THEN statement
 522 in this skolemised store gives $x_0 = y$. Since $x_0 = y$ is
 523 not a contradiction, the conditioner using the naïve ap-
 524 proach would be forced to keep in the whole IF–THEN
 525 statement, however powerful the theorem prover.

526 Using the less crude approach gives the two symbolic
 527 stores $[x \rightarrow y + 1]$ and $[x \rightarrow y + 2]$. The fact that in the
 528 loop, x is assigned an expression that is unaffected by the
 529 body of the loop has been taken into account. Since
 530 $y + 1 = y$ and $y + 2 = y$ are both contradictions,
 531 the IF statement following the WHILE loop can be
 532 removed.

533 4.2. The ConsSUS algorithm in detail

534 In this section, the algorithm used by ConsSUS is ex-
 535 plained in detail. For each WSL syntactic category, the
 536 result of applying ConsSUS to it will be defined. It will be
 537 assumed that the starting conditioned state in each case
 538 is given by

$$540 \Sigma = \bigcup_{i=1}^n \{(b_i, \sigma_i)\}$$

541 where the b_i are boolean expressions representing path
 542 conditions and the σ_i are symbolic stores.

543 For each statement s , ConsSUS returns two objects:

- 544 • $state(\Sigma, s)$. The resulting conditioned state when con-
 545 ditioning statement s in Σ and
- 546 • $statement(\Sigma, s)$. The resulting simplified statement
 547 when conditioning statement s in Σ . If statement s
 548 is to be removed by ConsSUS, it returns SKIP. A final
 549 post-processing phase will call *FerMat*'s Delete-
 550 All_Skips transformation to remove all the
 551 SKIPS that have introduced by performing this
 552 operation.

553 Calls to the theorem prover, *FerMat* Simplify will
 554 be represented by the expression $prove(b)$, where b is a
 555 boolean expression. The expression, $prove(b)$, is defined
 556 to return *true* if the theorem prover determines that b is
 557 valid and *false* otherwise. If $prove(b)$ returns *false*, this
 558 represents the fact that *either* the theorem prover cannot
 559 reduce the condition to *true* or it reduces it to the condi-
 560 tion to *false*. Given a conditioned state, Σ , and a boo-
 561 lean expression, b , we define $AllImply(\Sigma, b)$ to be *true*
 562 if and only if, for all pairs (c, σ) in conditioned state Σ ,
 563 $prove(c \Rightarrow \sigma b)$ evaluates to *true*. Where, given a symbolic
 564 store σ , the expression, σb , denotes the result of symbol-
 565 ically evaluating b in σ .

Suppose b is the guard of an IF statement. $AllIm-$
 $ply(\Sigma, b)$ implies that the THEN branch *must* be executed
 in Σ and the ELSE branch can be removed. Similarly
 $AllImply(\Sigma, \text{NOT } b)$ implies that THEN branch can be re-
 moved. Suppose b is a guard of a WHILE loop, then $All-$
 $Imply(\Sigma, b)$ implies that the body of the loop is executed
 at least once and $AllImply(\Sigma, \text{NOT } b)$ implies that the
 loop body is not executed at all.

4.2.1. Conditioning ABORT

In order to condition an ABORT statement, a special
 conditioned state called the ABORT state is introduced
 and written \perp . It consists of the single pair $(false, id)$.

$$state(\Sigma, \text{ABORT}) \triangleq \perp$$

$$statement(\Sigma, \text{ABORT}) \triangleq \text{ABORT}$$

For all statements s , define

$$state(\perp, s) \triangleq \perp$$

$$statement(\perp, s) \triangleq \text{SKIP}$$

This guarantees that all statements following an ABORT
 will be removed. In the rest of the discussion it is as-
 sumed that $\Sigma \neq \perp$.

4.2.2. Conditioning SKIP

$$state(\Sigma, \text{SKIP}) \triangleq \Sigma$$

$$statement(\Sigma, \text{SKIP}) \triangleq \text{SKIP}$$

Conditioning a SKIP has no effect.

4.2.3. Conditioning assert statements

In WSL, an assert statement is written $\{b\}$ where b is
 a boolean expression. It is semantically equivalent to
 IF b THEN SKIP ELSE ABORT FI. There are three
 cases to consider:

Case	Condition	Meaning
1	$AllImply(\Sigma, b)$	The assert condition will always be <i>true</i>
2	$AllImply(\Sigma, \text{NOT } b)$	The assert condition will always be <i>false</i>
3	None of the above	Nothing can be inferred

From the semantics of the Assert statement it is clear
 that in case 1, the Assert is equivalent to SKIP so the
 rules for SKIP above apply. In case 2, the Assert is
 equivalent to ABORT so the rules for ABORT above ap-
 ply. If neither the guard of the Assert is not always *true*
 or not always *false* in the current state, then the Assert
 cannot be removed. The resulting state will have the

604 same set of symbolic stores. The path conditions of the
605 resulting state will be different however. For each pair,
606 (b_i, σ_i) the resulting state will have a corresponding pair
607 $(b_i \text{ AND } \sigma_i b, \sigma_i)$ where $b_i \text{ AND } \sigma_i b$ is the boolean expression
608 created by ‘ANDing’ the boolean expression b_i with the
609 result of symbolically evaluating the boolean expres-
610 sion ⁶ b in symbolic store σ_i .

611 This represents the fact that a program will continue
612 executing after an Assert statement in stores where b
613 evaluates to *true*. Formally, in this case

$$615 \text{ state}(\Sigma, \{b\}) \triangleq \bigcup_{i=1}^n \{(b_i \text{ AND } \sigma_i b, \sigma_i)\}.$$

$$617 \text{ statement}(\Sigma, \{b\}) \triangleq \{b\}.$$

618 4.2.4. Conditioning assignment statements

619 When conditioning assignment statements, *ConsUS*
620 symbolically evaluates the expression on the right-hand
621 side of the assignment and updates the symbolic stores
622 accordingly. The path conditions do not change. In or-
623 der to symbolically evaluate an expression e in a sym-
624 bolic store, σ , *ConsUS* replaces every variable in the
625 expression by its value in σ . Given a symbolic store, σ ,
626 we use standard notation $\sigma[x \rightarrow e]$ to represent a store
627 that ‘agrees’ with σ except that variable x is now mapped
628 to e . Using this, the conditioning of assignment state-
629 ments can be defined as follows:

$$\text{state}(\Sigma, x := e) \triangleq \bigcup_{i=1}^n \{(b, \sigma_i[x \rightarrow \sigma_i e])\}$$

$$631 \text{ statement}(\Sigma, x := e) \triangleq x := e.$$

632 4.2.5. Conditioning statements sequences

633 In the case of standard semantics (Stoy, 1985), the
634 meaning of a sequence of statements is the composition
635 of the meaning functions of the individual statements.
636 The same is true when conditioning:

$$\text{state}(\Sigma, s_1; s_2) \triangleq \text{state}(\text{state}(\Sigma, s_1), s_2)$$

$$\text{statement}(\Sigma, s_1; s_2) \triangleq \text{statement}(\Sigma, s_1);$$

$$638 \text{ statement}(\text{state}(\Sigma, s_1), s_2)$$

639 This reflects the fact that conditioned states are ‘passed
640 through’ the program in the same order that the pro-
641 gram would have been executed. Once again, if as a re-
642 sult of conditioning, both parts of the sequence reduce
643 to SKIP then they will both be removed by the post-
644 processing phase.

⁶ For example, if σ_i maps y to $z + 1$ and x to 17 and if b is the boolean expression: $y > x + 1$ and if b_i is the boolean expression: $a + z = 5$ then $(b_i \text{ AND } \sigma_i b)$ is the boolean expression: $a + z = 5 \text{ AND } z + 1 > 17 + 1$.

4.2.6. Conditioning guarded commands 645

646 In WSL, a generalised form of conditional known as
647 guarded command is used. A guarded command has
648 concrete syntax of the form 648

$$\text{IF } B_1 \text{ THEN } S_1 \text{ ELSEIF } \dots \text{ ELSEIF } B_n \text{ THEN } S_n \text{ FI.} \quad 650$$

651 Unlike the semantics of Dijkstra’s guarded commands
652 (Dijkstra, 1972), these are deterministic in the sense that
653 the guards are evaluated from left to right and when a
654 true one is found the corresponding body is executed.
655 If none of the guards evaluates to *true* then the program
656 aborts. Although WSL has conventional IF THEN ELSE
657 FI statement, these are implemented as a guarded com-
658 mand whose last guard is identically TRUE. An IF THEN
659 statement is also implemented as a guarded command
660 whose last guard is identically TRUE and whose corre-
661 sponding body is SKIP. For the purposes of describing
662 conditioning guarded commands, it is convenient to rep-
663 resent a guarded command as

$$B_1 \rightarrow S_1 \mid \dots \mid B_n \rightarrow S_n. \quad 665$$

666 Using WSL terminology, each $B_i \rightarrow S_i$ is known as a
667 guarded. Conditioning a guarded command is defined
668 in terms of conditioning a guarded, $B \rightarrow S$ so that is de-
669 fined first.

670 When conditioning a guarded, like in the case of the
671 Assert statement, there are three possibilities: 671

Case	Condition	Meaning
1	$AllImply(\Sigma, B)$	The guard B will always be <i>true</i>
2	$AllImply(\Sigma, \text{NOT } B)$	The guard B will always be <i>false</i>
3	None of the above	Nothing can be inferred

In cases 1 and 3 672

$$\text{state}(\Sigma, B \rightarrow S) \triangleq \text{state}(\Sigma', S)$$

$$\text{statement}(\Sigma, B \rightarrow S) \triangleq B \rightarrow \text{statement}(\Sigma', S) \quad 674$$

where 675

$$\Sigma' = \bigcup_{i=1}^n \{(b_i \text{ AND } \sigma_i B, \sigma_i)\}. \quad 677$$

678 In case 2, the guarded can be removed and the resulting
679 state will simply be Σ :

$$\text{state}(\Sigma, B \rightarrow S) \triangleq \Sigma$$

$$\text{statement}(\Sigma, B \rightarrow S) \triangleq \text{SKIP} \quad 681$$

682 Having defined how *ConsUS* conditions a single
683 guarded, we now return to define how *ConsUS* condi-
684 tions a complete guarded command. As already ex-
685 plained, a guarded command is a sequence of guarded: 685

687 $B_1 \rightarrow S_1 \mid \dots \mid B_n \rightarrow S_n$.

688 When conditioning a guarded command in Σ , the guard-
689 eds are conditioned, as described above, from left to
690 right. The j th guarded is conditioned in conditioned
691 state Σ_j where

693 $\Sigma_1 = \Sigma$

694 and

$$696 \quad \Sigma_{j+1} = \bigcup_{(b_i, \sigma_i) \in \Sigma_j} \{(b_i \text{ AND } \sigma_i B_j, \sigma_i)\}.$$

697 For each guarded, $B_j \rightarrow S_j$, *ConsUS* decides

- 698 (a) Whether to keep or remove it.
699 (b) Whether to continue processing the next guarded in
700 this guarded command or to move on to the next
701 statement after the guarded command.

702 Conditioning proceeds as follows:
703

- 704 • If *AllImPLY*(Σ_j, B_j) this implies that the j th guard will
705 be chosen in all paths where the previous guards have
706 not been chosen. The resulting statement will be
707 *statement*($\Sigma_j, B_j \rightarrow S_j$). Conditioning of the guarded
708 command can stop at this point since none of the
709 guarded to the right of this one will ever be executed
710 in Σ .
711 • If *AllImPLY*($\Sigma_j, \text{NOT } B_j$) this implies that the j th guard
712 will never be chosen. This guarded can, therefore,
713 be removed without conditioning it, and processing
714 can continue with the conditioning of the next
715 guarded, $B_{j+1} \rightarrow S_{j+1}$ in conditioned state $\Sigma_{j+1} = \Sigma_j$.
716 • If neither *AllImPLY*(Σ_j, B_j) nor *AllImPLY*($\Sigma_j, \text{NOT } B_j$)
717 then it cannot said for certain whether B_j will be cho-
718 sen or not. This is represented by keeping the
719 guarded, *statement*($\Sigma_j, B \rightarrow S_j$), and again moving
720 on to process the next guarded in conditioned state
721 Σ_{j+1} .
722

723 Processing continues in this way from left to right un-
724 til there are no more guarded to consider. The resulting
725 final conditioned state of the guarded command is the
726 union of all the conditioned states of the guarded that
727 were processed. The resulting final statement of the
728 guarded command is either:

- 729 (1) A guarded command consisting of the guarded
730 that were kept in by the above process, in the same
731 order (This rule only applies if more than one
732 guarded was kept in by the above process.) or
733 (2) The body of the only guarded that was kept in.
734 (This rule only applies if exactly one guarded was
735 kept in by the above process.) or
736 (3) ABORT (This rule only applies if no guarded were
737 kept in by the above process.)

Since, as described above, not all guarded need neces- 738
sarily be processed, this algorithm is, in effect, pruning 739
infeasible paths ‘on the fly’. This is a much more efficient 740
approach than that of *ConsIT*, where all paths were 741
fully expanded before any simplification took place. 742

4.2.7. Conditioning loops 743

Before the result of conditioning **WHILE B DO S OD**, 744
in conditioned state Σ is defined, some preliminary def- 745
initions are required. 746

Definition 1. Σ^{true} is the initial state Σ with the added 747
constraint that the guard, B , is initially true in all pairs 748
of Σ . 749

$$\Sigma^{true} = \bigcup_{(b, \sigma) \in \Sigma} \{(b \text{ AND } (\sigma B), \sigma)\}.$$

Similarly, 752

Definition 2. Σ^{false} is the initial state Σ with the added 753
constraint that the guard, B , is initially false in all pairs 754
of Σ . 755

$$\Sigma^{false} = \bigcup_{(b, \sigma) \in \Sigma} \{(b \text{ AND } (\sigma \text{ NOT } B), \sigma)\}.$$

Definition 3. (The Skolemised Conditioned State, 758
 Σ') The skolemised conditioned state

$$\Sigma' = \bigcup_{(b, \sigma) \in \Sigma^{true}} \{(b, \sigma')\}.$$

where the symbolic stores, σ'_i , are the skolemised ver- 762
sions of the σ_i with respect to S , as described in Section 763
4.1. 764

Definition 4. ($\Sigma^{\geq 1}$) $\Sigma^{\geq 1}$ is the conditioned state after 765
exactly one execution of loop in state Σ 766

$$\Sigma^{\geq 1} = \text{state}(\Sigma', S).$$

where the symbolic stores, σ'_i , are the skolemised ver- 769
sions of the σ_i with respect to S . 770

Definition 5. (Σ^{final}) Σ^{final} is the final conditioned state 771
after at least one execution of the loop in state Σ assum- 772
ing that the loop terminates 773

$$\Sigma^{final} = \bigcup_{(b, \sigma) \in \Sigma^{\geq 1}} \{(b \text{ AND } \sigma (\text{NOT } B), \sigma)\}.$$

When conditioning a loop of the form **WHILE B DO S** 776
OD, in conditioned state Σ , *ConsUS* checks all the seven 777
conditions in the table in Fig. 7. 778

Each case in Fig. 7 has the following implications. 779

Case 1	Loop not executed
Case 2	Nothing known
Case 3	If loop executed once, then it does not terminate
Case 4	If loop executed once, then it executes exactly once
Case 5	Loop executes at least once
Case 6	Loop non-terminates
Case 7	Loop executes exactly once

	Final State
Case 1 (Loop not executed)	Σ
Case 2 (Nothing known)	$\Sigma^{false} \cup \Sigma^{final}$
Case 3 (If once, non-termination)	Σ^{false}
Case 4 (If once, exactly once)	$state(\Sigma, IF B THEN S FI)$
Case 5 (At least once)	Σ^{final}
Case 6 (Non-termination)	\perp
Case 7 (Exactly once)	$state(\Sigma, S)$

Fig. 8. WHILE loop final states in each case.

780 Blank entries in the table mean we do not care about
 781 these values. The other combinations not considered are
 782 all impossible. For each of these cases

784 $state(\Sigma, WHILE B DO S OD)$

785 and

787 $statement(\Sigma, WHILE B DO S OD)$

788 will have different values (Figs. 8 and 9). Each is now
 789 considered in turn.

Case 1: the loop is not executed. There is no change to the final conditioned state and loop can be removed.

Case 2: nothing is known about the loop. The final conditioned state is the union of the final conditioned states corresponding to not executing the loop at all and to terminating after at least one execution. It is not necessary to consider non-termination as no states after non-termination are reachable. The resulting statement is the while loop with its body conditioned in Σ' where Σ' is the skolemised state.

Case 3: if the loop is executed at least once then it non-terminates. The final conditioned state corresponds to not executing the loop, since this is the only way termination can occur. The loop can be replaced with an assertion of the negation of the guard.

Case 4: if the loop is executed once then it executes at most once. This is equivalent to conditioning the corresponding conditional statement in state Σ .

	Final Statement
Case 1 (Loop not executed)	SKIP
Case 2 (Nothing known)	$WHILE B DO statement(\Sigma', S) OD$
Case 3 (If once, non-termination)	$\{ NOT B \}$
Case 4 (If once, exactly once)	$statement(\Sigma, IF B THEN S FI)$
Case 5 (At least once)	$WHILE B DO statement(\Sigma', S) OD$
Case 6 (Non-termination)	ABORT
Case 7 (Exactly once)	$statement(\Sigma, S)$

Fig. 9. WHILE loop resulting statements in each case.

Case 5: the loop is executed at least once. The final conditioned state is the Σ^{final} , corresponding to the loop terminating after at least one execution. It is not necessary to consider non-termination as no states after non-termination are reachable. The resulting statement is the while loop with its body conditioned in skolemised state, Σ' .

Case 6: the loop does not terminate. The final state is \perp and the loop can be replaced with ABORT.

Case 7: The loop executes exactly once. This is equivalent to conditioning S in Σ . Since $AllImply(\Sigma, B)$ and $AllImply(\Sigma^{\geq 1}, NOT B)$ we do not need to add the constraints that the loop guard is initially true and finally false.

	$AllImply(\Sigma, NOT B)$	$AllImply(\Sigma, B)$	$AllImply(\Sigma^{\geq 1}, NOT B)$	$AllImply(\Sigma^{\geq 1}, B)$
Case 1	T			
Case 2	F	F	F	F
Case 3	F	F	F	T
Case 4	F	F	T	F
Case 5	F	T	F	F
Case 6	F	T	F	T
Case 7	F	T	T	F

Fig. 7. WHILE loop possibilities.

823 4.3. Examples

824 This section gives examples of the output of *ConSUS*
 825 for a variety of small examples in order to demonstrate
 826 its behaviour

827 The program in Fig. 10 is an example with two con-
 828 secutive identical while loops. *ConSUS* removes the sec-
 829 ond loop since its guard can never be true after
 830 completing execution of the first loop. This is true even
 831 if the first loop is not executed or if it non-terminates.

832 In Fig. 11 there is a loop which if executed once never
 833 terminates. *ConSUS* replaces this loop with an Assert
 834 statement that asserts that the guard of the loop is false.
 835 *ConSUS* also recognised that to 'get past' the loop, it
 836 must not be executed and therefore the initial assign-
 837 ment to *x* is not overwritten and therefore the following
 838 IF statement can be simplified.

839 The program in Fig. 12 has a while loop which is
 840 either not executed at all or exactly once. *ConSUS* re-
 841 places it with an IF. In the current implementation, if
 842 the 2 was replaced by $x + 1$, say, no simplification would
 843 take place. This is because the *ConSUS* infers that only a
 844 single loop iteration is possible by analysing the loop
 845 guard in the skolemised state and not in the state after
 846 a single execution.

847 In Fig. 13, although the loop itself cannot be simpli-
 848 fied, *ConSUS* recognises that the loop must be executed
 849 at least once and hence the later IF can be simplified.

<pre> WHILE x<1 DO x:=x+1 OD; WHILE x<1 DO x:=x+1 OD </pre>	<pre> WHILE x<1 DO x:=x+1 OD </pre>
Original Program	Output from <i>ConSUS</i>

Fig. 10. Conditioning a WHILE loop (Case 1).

<pre> x:=p; WHILE x>0 DO x:=1 OD; IF x=p THEN y:=2 ELSE y:=1 FI </pre>	<pre> x:=p; {NOT x > 0}; y :=2 </pre>
Original Program	Output from <i>ConSUS</i>

Fig. 11. Conditioning a WHILE loop (Case 3).

<pre> WHILE x=1 DO x:=2 OD </pre>	<pre> IF x=1 THEN x:=2 FI </pre>
Original Program	Output from <i>ConSUS</i>

Fig. 12. Conditioning a WHILE loop (Case 4).

<pre> x:=1; WHILE x>0 DO x:=x+y; y:=2 OD; IF (y=2) THEN x:=1 ELSE x:=2 FI </pre>	<pre> x:=1; WHILE x>0 DO x:=x+y; y:=2 OD; x:=1 </pre>
Original Program	Output from <i>ConSUS</i>

Fig. 13. Conditioning a WHILE loop (Case 5).

In Fig. 14, *ConSUS* recognises that the program does not terminate and therefore everything apart from the initial Assert can be discarded since these statements are not reachable.

In Fig. 15 we have 'helped' the theorem prover with some knowledge that in this loop, *x* will always be greater than zero. From this, *ConSUS* has inferred that the loop will terminate after exactly one execution. As the implementation stands, without this human intervention, *ConSUS* would not produce simplification. As in

<pre> {x>1}; WHILE x>0 DO y:=x+y; OD; IF (x>0) THEN x:=1 ELSE x:=2 FI </pre>	<pre> {x>1}; </pre>
Original Program	Output from <i>ConSUS</i>

Fig. 14. Conditioning a WHILE loop (Case 6).

<pre> x:=1; WHILE x=1 DO x:=x+1; {x>1} OD; IF (x=1) THEN x:=1 ELSE x:=2 FI </pre>	<pre> x:=1 x:=x+1; {x>1} x:=2 </pre>
Original Program	Output from <i>ConsUS</i>

Fig. 15. Conditioning a WHILE loop (Case 7).

860 Case 4, this is because the *ConsUS* infers that only a sin-
861 gle loop iteration is possible by analysing the loop guard
862 in the skolemised state and not in the state after a single
863 execution. The algorithm could very straightforwardly
864 be changed to consider one iteration of the loop as a spe-
865 cial case. In this example, we see that slicing on x at the
866 end of the program before conditioning yields no simpli-
867 fication. But after conditioning, slicing on x at the end of
868 the program gives us the single statement $x := 2$.

869 5. Empirical validation

870 There are two sections to the empirical validation.
871 First, measurements of the effectiveness of *ConsUS*
872 when applied to small but realistic programs are per-
873 formed. The reduction in program size produced both
874 by conditioning alone and also produced by combining
875 conditioning with slicing are given. It is demonstrated
876 that in both cases conditioning using *ConsUS* produces
877 a considerable reduction in program size.

878 Secondly, the scalability of *ConsUS* is examined. Pro-
879 grams are constructed using multiple repetitions of one
880 of some fixed program fragments. This enables arbitrar-
881 ily large programs to be generated using these fragments.
882 *ConsUS* is timed on successive instances of various com-
883 binations of these fragments. Graphs of *ConsUS* execu-
884 tion time against program size are produced and
885 analysed. It is shown that the technique employed by
886 *ConsUS* appears to scale well at least at the unit level.

887 5.1. Effectiveness

888 In this section, the behaviour of *ConsUS* when ap-
889 plied to more realistic applications written in WSL is
890 analysed. Examples of the decrease in program size
891 when *ConsUS* is used on its own and also when it is
892 combined with a slicer are given for a variety of slicing
893 criteria. The programs⁷ considered are:

- (1) Student Marks Processor 894
- (2) A Calendar Program 895
- (3) Tax Allowance Calculator 896

897
898 As with other work on empirical aspects of slicing
899 (Harrold and Ci, 1998; Liang and Harrold, 1999; Mock
900 et al., 2002; Nishimatsu et al., 1999), the slicing criteria
901 are developed to be realistic criteria for the programs se-
902 lected, typical of the kinds of query a user of a slicing
903 system might present for the programs under
904 consideration.

5.1.1. Student marks processor 905

906 The student marks processor is a 230 line WSL pro-
907 gram which allows the user to enter marks for three
908 courses for an arbitrary number of students. It outputs
909 a variety of statistics including the final degree classifica-
910 tion for each student, the average, highest and lowest
911 marks for each course and the number of students
912 receiving each degree classification. The program has er-
913 ror checking in the sense that it uses loops to force the
914 user to input marks in the correct range (0–100). 914

915 The program is ‘decorated’ with a number of AS-
916 SERT statements of the form $\{\text{read} \geq 0 \text{ AND}$
917 $\text{read} \leq 100\}$. These correspond to the assumption
918 that all inputs to the program are correct at the first at-
919 tempt. When conditioned with respect to these criteria
920 *ConsUS* reduces the original 230 line program to 132
921 lines. All the unnecessary loops that force input to be
922 in the correct range are correctly removed by *ConsUS*.
923 This corresponds to a reduction in size as a result of con-
924 ditioning of almost 43%. 924

925 In the next set of experiments, the original program is
926 first sliced with respect to a number of different criteria
927 and then each of these slices is further conditioned. This
928 reduction in size produced by conditioning after slicing
929 is then recorded.

930 For example, the original program, was first sliced at
931 the end with respect to the variable `numfirsts` to give
932 rise to a 31 line program. The resulting slice was then
933 conditioned with respect to the same conditioning crite-
934 rion. Again, all the unnecessary loops that force input to
935 be in the correct range were removed by *ConsUS*. The
936 resulting conditioned slice, thus, has the same effect as
937 the original on the variable `numfirsts` assuming the
938 all inputs are correct first time. In this example there
939 was almost a 1/3 reduction in the size of the slice as a
940 result of conditioning after slicing.

941 Similarly, slicing with respect to variable `average3`
942 leaves a 29 line program which is further reduced to 19
943 by conditioning. Fig. 16 shows the results of these
944 experiments.

5.1.2. The calendar program 945

946 The input to this 123 line WSL program is any day
947 since first of January, 1 A.D. The output is the number 947

⁷ All the examples can be found at <http://www.doc.gold.ac.uk/~mas01sd/consus/>.

Conditioning Criterion	Slice Variable	Size of Original	Size after slicing	Size after conditioning	Reduction due to <i>ConSUS</i>
Inputs correct	No slicing	230 lines	230 lines	123 lines	43%
Inputs correct	numfirsts	230 lines	31 lines	21 lines	32%
Inputs correct	average3	230 lines	29 lines	19 lines	34%
Inputs correct	numstudents	230 lines	5 lines	3 lines	40%
Inputs correct	averagegrade1	230 lines	41 lines	32 lines	22%
Inputs correct	lowest1	230 lines	39 lines	24 lines	38%

Fig. 16. Size of Student Marks program after first slicing and then conditioning.

Conditioning Criterion	Slice Variable	Size of Original	Size after slicing	Size after conditioning	Reduction due to <i>ConSUS</i>
1	No Slicing	123 lines	123 lines	100 lines	19%
2	No Slicing	123 lines	123 lines	95 lines	23%
1	leap	123 lines	36 lines	32 lines	11%
2	leap	123 lines	36 lines	32 lines	11%
1	dayofweek	123 lines	90 lines	79 lines	12%
2	dayofweek	123 lines	90 lines	78 lines	13%
1	month	123 lines	21 lines	13 lines	38%
2	month	123 lines	21 lines	13 lines	38%

- Conditioning Criterion 1 corresponds to all inputs being initially in the correct range.
- Conditioning Criterion 2 corresponds to all inputs being initially in the correct range and the input year being after 1752.

Fig. 17. Size of Calendar program after first slicing and then conditioning.

948 of days since the first of January 1 A.D. together with
 949 the day of the week of the input date. The program takes
 950 account of the 11 'lost days' between 2 and 14 September
 951 1752 (*Calendar Act, 1751*) and the new rules for decid-
 952 ing on leap years that came into effect around 1800.

953 Again, measurements of the decrease in program size
 954 when *ConSUS* is used on its own and also when it is
 955 combined with a slicer were made.

956 The conditioning criteria are:

- 957 (1) that all inputs are first time correct (as in the Stu-
 958 dent Marks Processor Program) and
 959 (2) that all inputs are first time correct and that the
 960 year is after the dreaded 1752.

961 and the slices were taken at the end of the program with
 962 respect to variables:

- 963 (1) leap
 964 (2) dayofweek
 965 (3) month

966 Without first slicing, using the first conditioning crite-
 967 rion, *ConSUS* successfully removes all the 'force input
 968 loops' as before, reducing the program from 123 to
 969 100 lines and using the second conditioning criterion,

ConSUS also removes the code that specifically checks
 whether the date falls within the 'lost eleven days',
 reducing the original from 123 to 95 lines.

Slicing the program with respect to the variables
 above and then conditioning produces a further reduc-
 tion in program size as can be seen in Fig. 17.

5.1.3. The tax allowance calculator

The Tax Allowance Calculator is a 129 line program
 which asks the user a number of questions from which it
 then calculates the income tax allowance, tax code etc.
 for the user. It is an encoding of the UK Tax rules as
 they were between April 1998 and April 1999. Tax codes
 and allowances vary depending on the state of an indi-
 vidual. Important criteria include marital status, age
 and sex. A blind person gets a different allowance from
 a sighted person. This program is first sliced with respect
 to variables

- (1) code
 (2) pc10
 (3) personal

Each slice was then conditioned with respect to the
 criteria:

970
 971
 972
 973
 974
 975
 976
 977
 978
 979
 980
 981
 982
 983
 984
 985
 986
 987
 988
 989
 990
 991
 992

Conditioning Criterion	Slice Variable	Size of Original	Size after slicing	Size after conditioning	Reduction due to ConSUS
1	No Slicing	129 lines	129 lines	71 lines	45 %
2	No Slicing	129 lines	129 lines	71 lines	45 %
3	No Slicing	129 lines	129 lines	69 lines	47 %
1	code	129 lines	48 lines	24 lines	50 %
2	code	129 lines	48 lines	18 lines	63 %
3	code	129 lines	48 lines	24 lines	50 %
1	pc10	129 lines	42 lines	27 lines	36 %
2	pc10	129 lines	42 lines	26 lines	38 %
3	pc10	129 lines	42 lines	19 lines	55%
1	personal	129 lines	37 lines	27 lines	27%
2	personal	129 lines	37 lines	26 lines	30%
3	personal	129 lines	37 lines	24 lines	35%

- Conditioning Criterion 1 corresponds to the input age being greater than 65.
- Conditioning Criterion 2 corresponds to the input person being blind.
- Conditioning Criterion 3 corresponds to the input person being not married.

Fig. 18. Size of the Tax program after first slicing and then conditioning.

993 (1) age > 65
 994 (2) blind = 1
 995 (3) married = 0

996
 997 As in the other examples, Fig. 18 shows a significant
 998 reduction in program size results both from just condi-
 999 tioning and also from conditioning after slicing.

1000 5.1.4. Summary

1001 Conditioning the ‘small but realistic’ programs in our
 1002 study using ConSUS resulted in an average reduction in
 1003 program size of approximately 35%. The maximum
 1004 reduction was 63% and the minimum reduction was
 1005 11%. The average reduction in non-sliced programs
 1006 was 37% whereas the average for sliced programs was
 1007 34%. The conditioning criteria were not arbitrary but
 1008 represented ‘meaningful’ properties of the input state

of the programs. These figures show that conditioning 1009
 can produce a considerable reduction in program size 1010
 when applied to realistic programs. Clearly more work 1011
 is required to scale the application of conditioning to 1012
 larger programs. However, it should be stressed that 1013
 conditioning is inherently harder than traditional static 1014
 slicing due to the requirement of symbolic execution 1015
 and theorem proving and so the authors believe that 1016
 these initial results on small programs are encouraging, 1017
 going some way to demonstrating the ‘proof of concept’ 1018
 for conditioned slicing. 1019

5.2. Scalability 1020

Six classes of programs called F, T, SN, NSN, SSC, 1021
 and NSSC are considered. The programs in each class 1022
 are formed from ‘generative’ program fragments with 1023

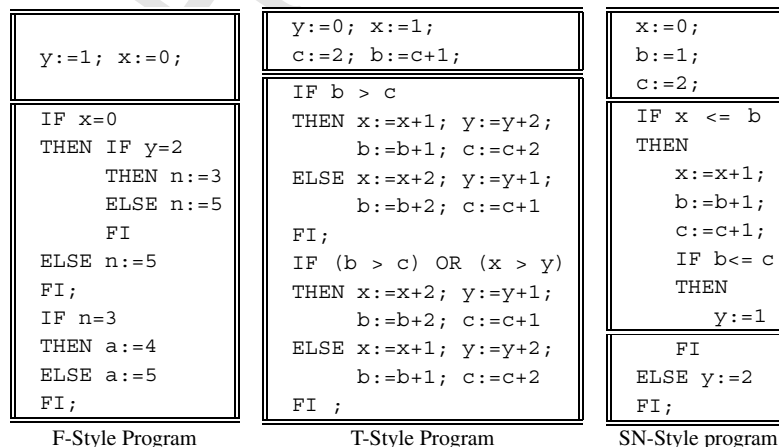


Fig. 19. The three considered classes of program.

```

{ (x>y) AND (y>z) };
IF (x>y) AND (y>z)
THEN IF (x>z)
    THEN x:=x+1;
        y:=y+1;
        z:=z+1
    ELSE a:=2
    FI
ELSE IF (x<=y) OR (y<=z)
    THEN x:=x+1;
        y:=y+1;
        z:=z+1
    ELSE a:=4
    FI
FI;

```

Fig. 20. Fragment for generating SSC and NSSC classes.

multiple repetitions of one of these fragments. The classes F, T, SN, NSN are generated from the 'base' programs given in Fig. 19. The SSC and NSSC are generated from fragment given in Fig. 20. This gives us a systematic approach to testing the scalability of *ConSUS*.

The programs of class F are generated from the fragments shown in Fig. 19 with multiple repetitions of the second fragment.

This set of programs tests the conditioning process of *ConSUS* on sequential IF statements where the predicates are testing equality of arithmetic expressions (as opposed to inequalities). Here the paths through the repetitions of the second fragment are always the same.

The T-class of programs is generated in the same manner using the fragments in Fig. 19, again repeating the second fragment. The conditions of the IF state-

ments involve inequalities and a logical OR. Furthermore, this class of programs involves greater symbolic evaluation than the F-class, as the program variables get updated continually (for example, $b := b + 1$) whereas in the F-class, the variables are assigned constant numeric values (for example, $n := 5$). Here the paths through the repetitions of the second fragment alternate for each repetition; with $b > c$ true and $(b > c)$ OR $(x > y)$ false first, and then vice-versa.


The SN-class of programs are generated from the SN-Style program in Fig. 19 by inserting multiple copies of the middle program fragment into the THEN branch of the previous IF statement, and adding an appropriate number of FIs in the third fragment. This produces an arbitrarily large nesting of IF statements.

The NSN-class is generated from the SN-Style program in Fig. 19 in exactly the same way except the initial fragment is excluded. The difference between these two program classes is that, in NSN, no simplification is possible using any conditioning process, whereas, in SN class of programs, the path through the program is uniquely determined.

The SSC-class of programs are generated from the fragments shown in Fig. 20 with multiple repetitions of the second fragment. The NSSC class is formed in exactly the same way except the initial fragment is excluded. The results of running *ConSUS* on a set of programs from each class are shown in Figs. 21–26. These results were obtained on a Dual Pentium III with 2×330 MHz and 512 MB RAM running Linux. The graphs show the time taken in seconds by *ConSUS* to condition a program of a given class, plotted against the size of the program in lines of code.

Least squares regression was performed on the data sets for the following models:

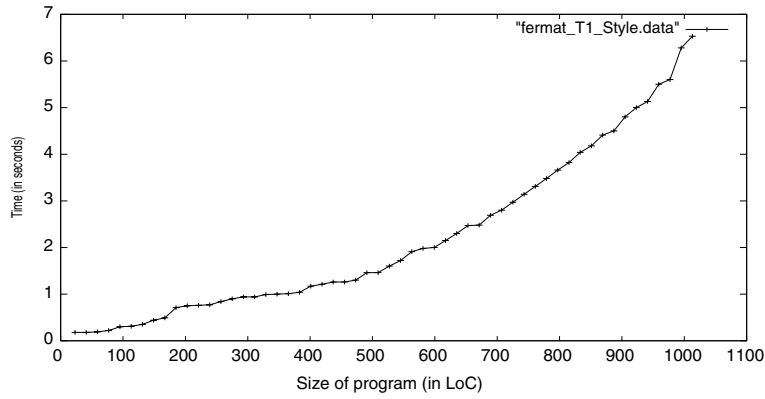
- linear model $y = a + bx$;
- exponential model $y = ae^{bx}$;



Using *FermaT*– Least squares quadratic polynomial:


$$y = 6.5956 \times 10^{-1} - 4.8090 \times 10^{-3}x + 4.0246 \times 10^{-5}x^2 \text{ with } R^2 = 0.99422$$

Fig. 21. Performance for F-class programs.



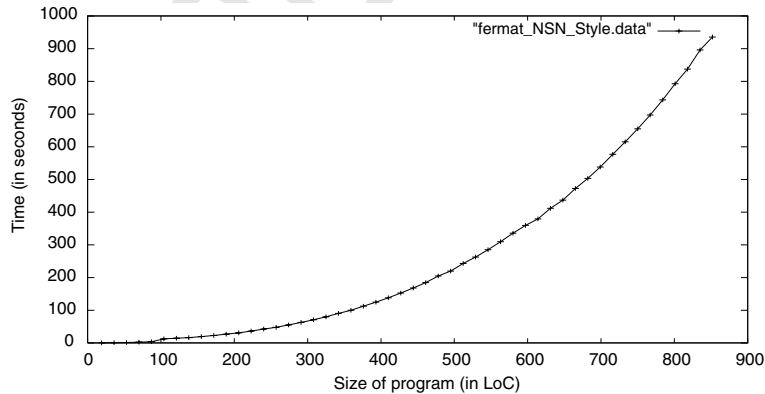
Using *FermaT* – Least squares quadratic polynomial:
 $y = 4.7372 \times 10^{-1} - 1.2072 \times 10^{-3}x + 6.6462 \times 10^{-6}x^2$ with $R^2 = 0.99155$.

Fig. 22. Performance for T-class programs.



Using *FermaT* – Least squares quadratic polynomial:
 $y = 4.2429 \times 10^1 - 4.1438 \times 10^{-1}x + 8.7712 \times 10^{-4}x^2$ with $R^2 = 0.98129$.

Fig. 23. Performance for SN-class programs.



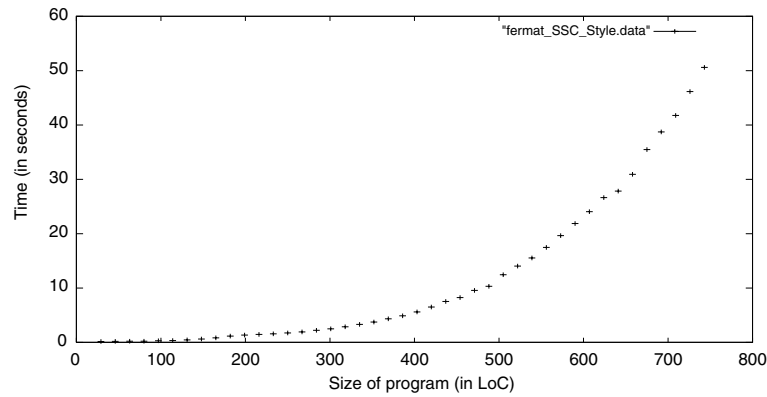
Using *FermaT* – Least squares quadratic polynomial:
 $y = 4.3645 \times 10^1 - 5.0660 \times 10^{-1}x + 1.7750 \times 10^{-3}x^2$ with $R^2 = 0.99652$.

Fig. 24. Performance for NSN-class programs.

- 1078 • power law model $y = ax^b$;
- 1079 • quadratic model $y = a + bx + cx^2$.

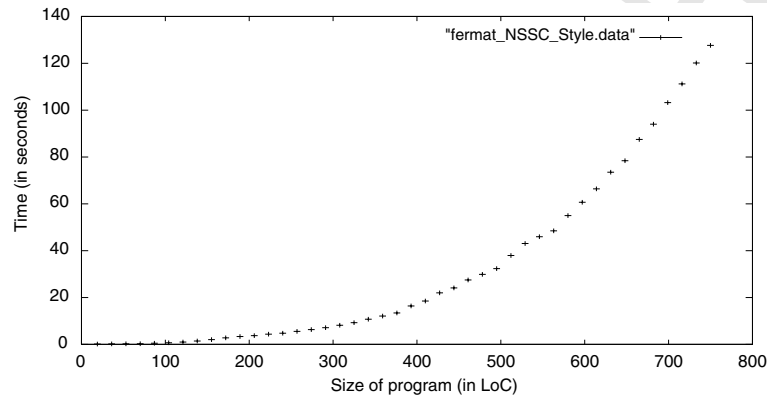
The quadratic model (with two degrees of freedom) gave the best fit to the data. The other models were signif-

1080
1081
1082



Using *FerMaT* – Least squares quadratic polynomial:
 $y = 15.9617 - 0.1498x + 2.79 \times 10^{-4}x^2$ with $R^2 = 0.977$.

Fig. 25. Performance for SSC-class programs.



Using *FerMaT* –Least squares quadratic polynomial:
 $y = 15.6851 - 0.1811x + 4.48 \times 10^{-4}x^2$ with $R^2 = 0.988$.

Fig. 26. Performance for NSSC-class programs.

1083 icantly worse even for models of one degree of freedom.
 1084 The least squares quadratic polynomials are given below
 1085 each figure along with the coefficient of determination R^2 .
 1086 Conditioning and conditioned slicing are typically
 1087 applied to programs at the unit level, for example, as a
 1088 support for detailed understanding (De Lucia et al.,
 1089 1996), as a unit level testing aid (Hierons et al., 2002)
 1090 or as a unit level reuse and code extraction tool (Cimitile
 1091 et al., 1995b; Cimitile et al., 1995a; Canfora et al.,
 1092 1994a). For these applications, quadratic performance
 1093 is acceptable and the technique therefore appears to
 1094 scale well, at least at the unit level.

1095 **6. Conclusions**

1096 The main contributions of this paper are:

- 1097 (1) To define a new more efficient algorithm and imple-
 1098 mentation for program conditioning which uses on-
 1099 the-fly pruning of symbolic execution paths.

- (2) To report on empirical studies which demonstrate that
 (a) On small ‘real programs’ this algorithm produces a considerable reduction in program size when used with and without a program slicer.
 (b) The *ConSUS* algorithm when used in conjunction with WSL’s *FerMaT Simplify* has the potential for ‘scaling up’ for use on larger systems.
 The algorithm defined in this paper is at the heart of *ConSUS*, a light-weight program conditioner for WSL. *ConSUS* prunes symbolic execution paths based on the validity of path conditions, thereby removing unreachable code. Unlike previous approaches, the *ConSUS* system integrates the reasoning and symbolic execution within a single system. The symbolic executor can eliminate paths which can be determined to be unexecutable in the current symbolic state. This pruning effect makes the algorithm more efficient. Furthermore, the reasoning is achieved, not using theorem proving, but rather using the in-built expression simplifier of *FerMaT*. This is a

1122 super-lightweight approach that may be capable of scal-
1123 ing to large programs.

1124 In the worst case, there is no doubt that the time to
1125 perform ‘best possible’ program conditioning will be
1126 exponential in the size of program being conditioned.
1127 This fact is inherent to the problem of conditioning.
1128 As in the case of theorem provers, this worst case sce-
1129 nario does not imply that implementations of condition-
1130 ers are infeasible. The reasons for this are twofold:

- 1131 (1) In conditioning, *any* resulting reduction in program
1132 size represents progress. Even if the best possible
1133 results require exponential time, it is possible that
1134 significant reductions will be performed more
1135 quickly.
- 1136 (2) In many cases conditioning may be performed in
1137 low order polynomial or even quadratic time as
1138 suggested by the empirical study in this paper. In
1139 such cases conditioning may be applicable to unit
1140 level applications at least.

1141 A ‘budget’ system similar to that used by *FermaT*
1142 *Simplify* is envisaged whereby conditioning is per-
1143 formed relative to a budget. The budget is reduced as
1144 processing takes place. When the budget expires further
1145 conditioning ceases.

1146 Clearly more work is required to scale the application
1147 of conditioning to larger programs. However, it should
1148 be stressed that conditioning is inherently harder than
1149 traditional static slicing due to the requirement of sym-
1150 bolic execution and theorem proving and so the authors
1151 believe that these initial results on small programs are
1152 encouraging, going some way to demonstrating the
1153 ‘proof of concept’ for conditioned slicing.

1155 It is possible that the performance achieved using a
1156 more powerful theorem prover, in some cases may out-
1157 weigh the light-weight approach when combined with
1158 ‘pruning on the fly’. This is because powerful reasoning
1159 power may result in ‘early pruning’ which may be missed
1160 by a less powerful theorem prover. An example of a
1161 more powerful theorem prover is the Co-operating
1162 Validity Checker(CVC) (Stump et al., 2002), the succes-
1163 sor to the Stanford Validity Checker (SVC) (Barrett et
1164 al., 1996). CVC is a high performance system for check-
1165 ing the validity of formulæ in a relatively rich decidable
1166 logic. CVC is applicable to boolean expressions made
1167 from these atoms.

1168 Very recently, CVC has also been incorporated into
1169 *ConSUS*. Early results do, indeed, show that the pro-
1170 gram in Fig. 5 is simplified when using the *ConSUS*
1171 algorithm in conjunction with CVC in place of *FermaT*
1172 *Simplify* which fails to remove any statements since it
1173 is unaware of the transitivity of \geq . Fig. 28 illustrates the
1174 result of conditioning the second code fragment in Fig.
1175 20 using both *FermaT* and CVC. Fig. 27 illustrates the
1176 result of conditioning the code fragments in Fig. 20
1177 using both *FermaT* and CVC. Future work will investi-
1178 gate the resulting trade off between speed and precision.
1179 Future work will also investigate whether it is possible
1180 to harness more of the power of a theorem prover like
1181 CVC in conditioning programs, for example performing
1182 induction on loop invariants. Other areas of interest in-
1183 clude the development of heuristics that allow both hea-
1184 vy and light weight approaches to be combined within
1185 the same conditioner and applications of this approach
1186 to backward conditioning (Fox et al., 2001).

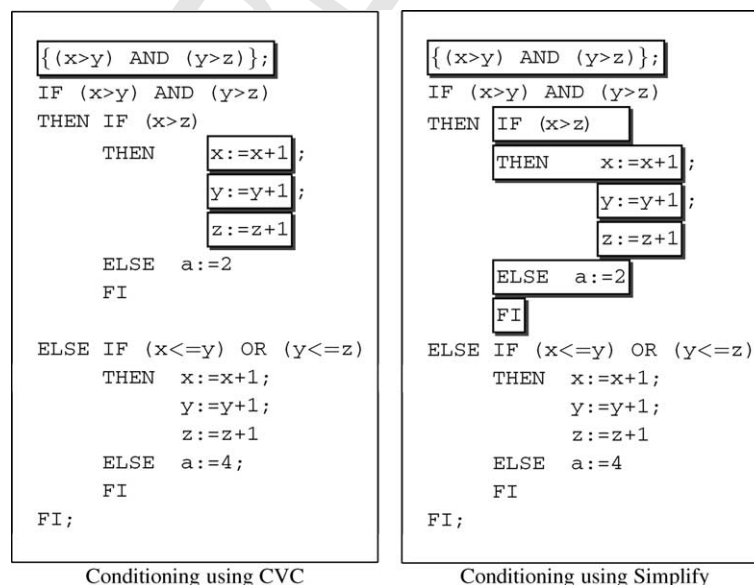


Fig. 27. Conditioning The SSC-Style program.

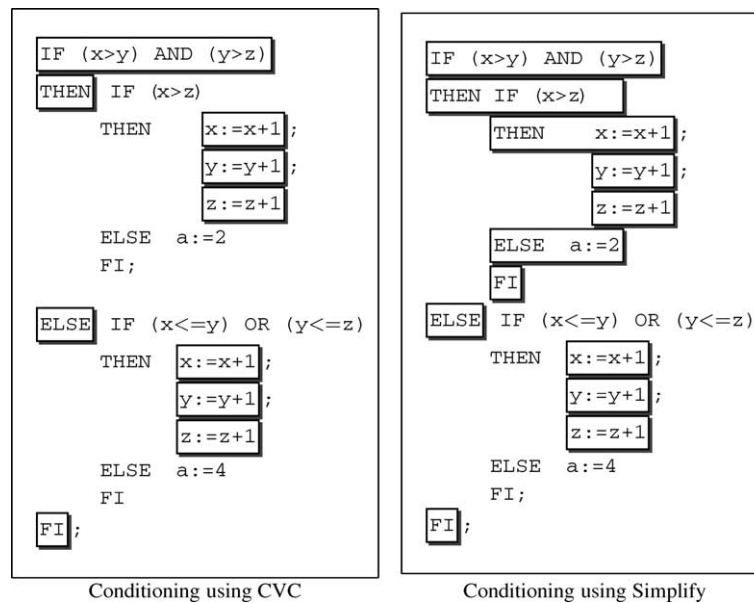


Fig. 28. Conditioning The NSSC-Style program.

1187 **References**

- 1188 Barrett, C., Dill, D., Levitt, J., 1996. Validity checking for combina- 1225
1189 tions of theories with equality. In: Srivas, M., Camilleri, A. (Eds.), 1226
1190 Formal methods in computer-aided design, Lecture Notes in 1227
1191 Computer Science, vol. 1166. Springer-Verlag, Berlin, pp. 187–201. 1228
1192 Bieman, J.M., Ott, L.M., 1994. Measuring functional cohesion. IEEE 1229
1193 Transactions on Software Engineering 20 (8), 644–657. 1230
1194 Binkley, D.W., 1998. The application of program slicing to regression 1231
1195 testing. In: Harman, M., Gallagher, K. (Eds.), Information and 1232
1196 Software Technology Special Issue on Program Slicing, 40. 1233
1197 Elsevier, Amsterdam, pp. 583–594. 1234
1198 Binkley, D.W., Gallagher, K.B., 1996. Program slicing. In: Zelkowitz, 1235
1199 M. (Ed.), Advances in Computing, vol. 43. Academic Press, New 1236
1200 York, pp. 1–50. 1237
1201 Binkley, D.W., Harman, M., to appear-a. A large-scale empirical study 1238
1202 of forward and backward static slice size and context sensitivity. In: 1239
1203 IEEE International Conference on Software Maintenance (ICSM 1240
1204 2003). IEEE Computer Society Press, Los Alamitos, CA, 1241
1205 Amsterdam. 1242
1206 Binkley, D.W., Harman, M., to appear-b. A survey of empirical results 1243
1207 on program slicing. Advances in Computers. 1244
1208 Binkley, D.W., Harman, M., Raszewski, L.R., Smith, C., 2000. An 1245
1209 empirical study of amorphous slicing as a program comprehension 1246
1210 support tool. In: 8th IEEE International Workshop on Program 1247
1211 Comprehension (IWPC 2000). IEEE Computer Society Press, Los 1248
1212 Alamitos, CA, USA, Limerick, Ireland, pp. 161–170. 1249
1213 Binkley, D.W., Horwitz, S., Reps, T., 1995. Program integration for 1250
1214 languages with procedure calls. ACM Transactions on Software 1251
1215 Engineering and Methodology 4 (1), 3–35. 1252
1216 Calendar Act, 1751. Calendar Act, Anno vicesimo quarto George II, 1253
1217 cap. xxiii. 1254
1218 Canfora, G., Cimitile, A., De Lucia, A., 1998. Conditioned program 1255
1219 slicing. In: Harman, M., Gallagher, K. (Eds.), Information and 1256
1220 Software Technology Special Issue on Program Slicing, vol. 40. 1257
1221 Elsevier Science, Amsterdam, pp. 595–607. 1258
1222 Canfora, G., Cimitile, A., De Lucia, A., Lucca, G.A.D., 1994a. 1259
1223 Software salvaging based on conditions. In: International Confer- 1260
1224 ence on Software Maintenance (ICSM'96). IEEE Computer 1261
1262
1263 Society Press, Los Alamitos, CA, USA, Victoria, Canada, pp. 424–433. 1264
1265 Canfora, G., Cimitile, A., Munro, M., 1994b. RE²: reverse engineering 1266
1267 and reuse re-engineering. Journal of Software Maintenance: 1267
1268 Research and Practice 6 (2), 53–72. 1268
1269 Cimitile, A., De Lucia, A., Munro, M., 1995a. Identifying reusable 1269
1270 functions using specification driven program slicing: a case study. 1270
1271 In: Proceedings of the IEEE International Conference on Software 1271
1272 Maintenance (ICSM'95). IEEE Computer Society Press, Los 1272
1273 Alamitos, CA, USA, Nice, France, pp. 124–133. 1273
1274 Cimitile, A., De Lucia, A., Munro, M., 1995b. Qualifying reusable 1274
1275 functions using symbolic execution. In: Proceedings of the 2nd 1275
1276 working conference on reverse engineering. IEEE Computer 1276
1277 Society Press, Los Alamitos, CA, USA, Toronto, Canada, pp. 1277
1278 178–187. 1278
1279 Coen-Portisini, A., De Paoli, K., 1990. SYMBAD: A symbolic executor 1279
1280 of sequential Ada programs. In: IFAC SAFECOMP'90. London, 1280
1281 pp. 105–111. 1281
1282 Coen-Portisini, A., De Paoli, R., Ghezzi, C., Mandrioli, D., 1991. 1282
1283 Software specialization via symbolic execution. IEEE Transactions 1283
1284 on Software Engineering 17 (9), 884–899. 1284
1285 Coward, P.D., 1988. Symbolic execution systems—a review. Software 1285
1286 Engineering Journal 3 (6), 229–239. 1286
1287 Coward, P.D., 1991. Symbolic execution and testing. Information and 1287
1288 Software Technology 33 (1), 53–64. 1288
1289 Danicic, S., Fox, C., Harman, M., Hierons, R.M., 2000. ConSIT: A 1289
1290 conditioned program slicer. In: IEEE International Conference on 1290
1291 Software Maintenance (ICSM'00). IEEE Computer Society Press, 1291
1292 Los Alamitos, CA, USA, pp. 216–226. 1292
1293 Danicic, S., Fox, C., Harman, M., Hierons, R.M., accepted. The 1293
1294 ConSIT conditioned slicing system. Software Practice and 1294
1295 Experience. 1295
1296 De Lucia, A., 2001. Program slicing: methods and applications. In: 1st 1296
1297 IEEE International Workshop on Source Code Analysis and 1297
1298 Manipulation. IEEE Computer Society Press, Los Alamitos, CA, 1298
1299 USA, Florence, Italy, pp. 142–149. 1299
1300 De Lucia, A., Fasolino, A.R., Munro, M., 1996. Understanding 1300
1301 function behaviours through program slicing. In: Proceedings of 1301
1302 the 4th IEEE Workshop on Program Comprehension. IEEE 1302
1303 1263

- 1264 Computer Society Press, Los Alamitos, CA, USA, Berlin, Ger-
 1265 many, pp. 9–18.
- 1266 DeMillo, R.A., Offutt, A.J., 1993. Experimental results from an
 1267 automatic test generator. *ACM Transactions of Software Engi-
 1268 neering and Methodology* 2 (2), 109–127.
- 1269 Dijkstra, E.W., 1972. *A discipline of programming*. Prentice-Hall,
 1270 Englewood Cliffs, NJ.
- 1271 Ershov, A.P., 1978. *On the essence of computation*. North-Holland,
 1272 Amsterdam, pp. 391–420.
- 1273 Field, J., Ramalingam, G., Tip, F., 1995. Parametric program slicing.
 1274 In: *22nd ACM Symposium on Principles of Programming Lan-
 1275 guages*. San Francisco, CA, pp. 379–392.
- 1276 Fox, C., Harman, M., Hierons, R.M., Danicic, S., 2001. Backward
 1277 conditioning: a new program specialisation technique and its
 1278 application to program comprehension. In: *Proceedings of the 9th
 1279 IEEE International Workshop on Program Comprehesion
 1280 (IWPC'01)*. Computer Society Press, Los Alamitos, CA, USA,
 1281 Toronto, Canada, pp. 89–97.
- 1282 Futamura, Y., 1971. Partial evaluation of computation process—an
 1283 approach to a compiler compiler. *Systems, Computers, Controls* 2
 1284 (5), 721–728.
- 1285 Gallagher, K.B., Lyle, J.R., 1991. Using program slicing in software
 1286 maintenance. *IEEE Transactions on Software Engineering* 17 (8),
 1287 751–761.
- 1288 Girgis, M.R., 1992. An experimental evaluation of a symbolic
 1289 execution system. *Software Engineering Journal* 7 (4), 285–290.
- 1290 Grammatech Inc., 2002. The codesurfer slicing system. Available from:
 1291 <<http://www.grammatech.com>>.
- 1292 Harman, M., Danicic, S., 1995. Using program slicing to simplify
 1293 testing. *Software Testing, Verification and Reliability* 5 (3), 143–
 1294 162.
- 1295 Harman, M., Hierons, R.M., 2001. An overview of program slicing.
 1296 *Software Focus* 2 (3), 85–92.
- 1297 Harman, M., Hierons, R.M., Danicic, S., Howroyd, J., Fox, C., 2001.
 1298 Pre/post conditioned slicing. In: *IEEE International Conference on
 1299 Software Maintenance (ICSM'01)*. IEEE Computer Society Press,
 1300 Los Alamitos, CA, USA, Florence, Italy, pp. 138–147.
- 1301 Harrold, M.J., Ci, N., 1998. Reuse-driven interprocedural slicing. In:
 1302 *Proceedings of the 20th International Conference on Software
 1303 Engineering*. IEEE Computer Society Press, Los Alamitos, CA,
 1304 USA, pp. 74–83.
- 1305 Hausler, P.A., 1989. Denotational program slicing. In: *Proceedings of
 1306 the 22nd Annual Hawaii International Conference on System
 1307 Sciences*, vol. II, pp. 486–495.
- 1308 Hierons, R.M., Harman, M., Danicic, S., 1999. Using program slicing
 1309 to assist in the detection of equivalent mutants. *Software Testing,
 1310 Verification and Reliability* 9 (4), 233–262.
- 1311 Hierons, R.M., Harman, M., Fox, C., Ouarbya, L., Daoudi, M., 2002.
 1312 Conditioned slicing supports partition testing. *Software Testing,
 1313 Verification and Reliability* 12, 23–28.
- 1314 Horwitz, S., Prins, J., Reps, T., 1989. Integrating non-interfering
 1315 versions of programs. *ACM Transactions on Programming Lan-
 1316 guages and Systems* 11 (3), 345–387.
- 1317 King, J.C., 1976. Symbolic execution and program testing. *Communi-
 1318 cations of the ACM* 19 (7), 385–394.
- 1319 King, K.N., Offutt, A.J., 1991. A FORTRAN language system for
 1320 mutation-based software testing. *Software Practice and Experience*
 1321 21, 686–718.
- 1322 Korel, B., Laski, J., 1988. Dynamic program slicing. *Information
 1323 Processing Letters* 29 (3), 155–163.
- 1324 Korel, B., Rilling, J., 1998. Dynamic program slicing methods. In:
 1325 Harman, M., Gallagher, K. (Eds.), *Information and Software
 1326 Technology Special Issue on Program Slicing*, vol. 40. Elsevier,
 1327 Amsterdam, pp. 647–659.
- 1328 Krinke, J., Snelting, G., 1998. Validation of measurement software as
 1329 an application of slicing and constraint solving. In: Harman, M.,
 1330 Gallagher, K. (Eds.), *Information and Software Technology
 Special Issue on Program Slicing*, vol. 40. Elsevier, Amsterdam,
 pp. 661–675.
- Liang, D., Harrold, M.J., 1999. Reuse-driven interprocedural slicing in
 the presence of pointers and recursion. In: *International Confer-
 ence of Software Maintenance*. IEEE Computer Society Press, Los
 Alamitos, CA, USA, Oxford, UK, pp. 410–430.
- Lyle, J.R., Weiser, M., 1987. Automatic program bug location by
 program slicing. In: *Proceedings of the 2nd International Confer-
 ence on Computers and Applications*. IEEE Computer Society
 Press, Los Alamitos, CA, USA, Peking, pp. 877–882.
- Mock, M., Atkinson, D.C., Chambers, C., Eggers, S.J., 2002. Improving
 program slicing with dynamic points-to data. In: Griswold, W.G. (Ed.),
*Proceedings of the 10th ACM SIGSOFT Symposium on the Founda-
 tions of Software Engineering (FSE-02)*. ACM Press, New York, pp.
 71–80.
- Moskewicz, M., Madigan, C., Zhao, Y., Zhang, L., Malik, S., 2001. Chaff:
 engineering an efficient SAT solver. In: *39th Design Automation
 Conference*, Las Vegas.
- Nishimatsu, A., Jihira, M., Kusumoto, S., Inoue, K., 1999. Call-mark
 slicing: an efficient and economical way of reducing slices. In:
*Proceedings of the 21st International Conference on Software
 Engineering*. ACM Press, New York, pp. 422–431.
- Offutt, A.J., 1990. An integrated system for automatically generating
 test data. In: Ng, P.A., Ramamoorthy, C.V., Seifert, L.C., Yeh,
 R.T. (Eds.), *Proceedings of the First International Conference on
 Systems Integration*. IEEE Computer Society Press, Morristown,
 NJ, pp. 694–701.
- Ott, L.M., Thuss, J.J., 1989. The relationship between slices and
 module cohesion. In: *Proceedings of the II ACM conference on
 Software Engineering*, pp. 198–204.
- Ouarbya, L., Danicic, S., Daoudi, D.M., Harman, M., Fox, C., 2002. A
 de-notational interprocedural program slicer. In: *IEEE Working
 Conference on Reverse Engineering (WCRE 2002)*. IEEE Compu-
 ter Society Press, Los Alamitos, CA, Richmond, VA, USA, pp.
 181–189.
- Stoy, J.E., 1985. *Denotational semantics: The Scott-Strachey approach
 to programming language theory*, third ed. MIT Press, Cambridge,
 MA.
- Stump, A., Barrett, C.W., Dill, D.L., 2002. CVC: a cooperating
 validity checker. In: Godskesen, J.C. (Ed.), *Proceedings of the
 International Conference on Computer-Aided Verification*. Lecture
 Notes in Computer Science.
- Tip, E., 1995. A survey of program slicing techniques. *Journal of
 Programming Languages* 3 (3), 121–189.
- Venkatesh, G.A., 1991. The semantic approach to program slicing. In:
*ACM SIGPLAN Conference on Programming Language Design
 and Implementation*. Toronto, Canada, pp. 26–28 (*Proceedings in
 SIGPLAN Notices* 26(6), pp. 107–119, 1991).
- Ward, M., 1989. *Proving program refinements and transformations*.
 DPhil Thesis, Oxford University.
- Ward, M., 1994. Reverse engineering through formal transformation.
The Computer Journal 37 (5), 795–813.
- Ward, M., 1999. Assembler to C migration using the FermaT
 transformation system. In: *IEEE International Conference on
 Software Maintenance (ICSM'99)*. IEEE Computer Society Press,
 Los Alamitos, CA, USA, Oxford, UK.
- Weiser, M., 1979. Program slices: Formal, psychological, and practical
 investigations of an automatic program abstraction method. PhD
 thesis, University of Michigan, Ann Arbor, MI.
- Weiser, M., 1982. Programmers use slicing when debugging. *Communi-
 cations of the ACM* 25 (7), 446–452.
- Weiser, M., 1984. Program slicing. *IEEE Transactions on Software
 Engineering* 10 (4), 352–357.
- Special Issue on Program Slicing, vol. 40. Elsevier, Amsterdam, pp. 661–675.
- 1331
1332
1333
1334
1335
1336
1337
1338
1339
1340
1341
1342
1343
1344
1345
1346
1347
1348
1349
1350
1351
1352
1353
1354
1355
1356
1357
1358
1359
1360
1361
1362
1363
1364
1365
1366
1367
1368
1369
1370
1371
1372
1373
1374
1375
1376
1377
1378
1379
1380
1381
1382
1383
1384
1385
1386
1387
1388
1389
1390
1391
1392
1393
1394
1395
1396
1397
1398

Sebastian Danicic is a lecturer in Computer Science at Goldsmiths College, University of London. He has a BSc in Pure Mathematics from Queen Mary College, Univeristy London, an MSc from Oxford

1399 University, and a PhD from the univeristy of North London. He has
1400 worked extensively on program slicing, transformation and other areas
1401 of theoretical computer science.

1402
1403 **Mohammed Daoudi** was admitted by Goldsmiths College to the Uni-
1405 versity of London degree of Bachelor of Science with First Class
1406 Honours in the field of Mathematics, in 2000. Since then, he has been a
1407 PhD student at Goldsmiths under the supervision of Sebastian
1408 Danicic and part of the Program Transformation and Analysis Group
1409 in the Department of Computing at the same college. I am also a
1410 member of the Vastt Research Group. His research interests include:
1411 program slicing and program conditioning. His work is concerned with
1412 theoretical and practical aspects of static program analysis. He is
1413 investigating a variety of partial semantic preserving program simpli-
1414 fication methods primarily using techniques associated with slicing. He
1415 has developed a system that performs program conditioning with the
1416 aid of various simplifiers. He has published work and is continuing to
1417 research in the area of Conditioned Slicing.

1418
1419 **Chris Fox** obtained his PhD in Computer Science from the University
1421 of Essex in 1993. His research interests are in software analysis,
1422 foundations of formal semantics, and business process modelling. He
1423 has held positions at the University of Saarland (Saarbrücken, Ger-
1424 many), Goldsmiths College (University of London), King's College
1425 London (University of London), and the University of Essex, where he
1426 is currently a Reader in Computer Science.

1427
1428 **Mark Harman** is a Reader in Computing at Brunel University, UK. He
1430 has worked extensively on program slicing, transformation and testing
1431 and more recently he founded the field of search-based software
1432 engineering. Dr. Harman is funded by the UK Engineering and
1433 Physical Sciences Research Council (EPSRC) and DaimlerChrysler,

Berlin. He is program chair for several conferences and workshops
including the 20th international conference on software maintenance
and the search-based software engineering track of the genetic and
evolutionary computation conference.

Rob Hierons received a BA in Mathematics (Trinity College, Cam-
bridge), and a PhD in Computer Science (Brunel University). He then
joined the Department of Mathematical and Computing Sciences at
Goldsmiths College, University of London, before returning to Brunel
University in 2000. He was promoted to Professor in 2003.

John Howroyd received BA in Mathematics from Oxford university in
1990, and a PhD in Mathematics from University College London in
1995 He has worked as a researcher in the University of St Andrews
1993–1995 and as a Lecturer in Mathematics at Goldsmiths College
London since 1996. His research interests include Geometric measure
theory, Functional Analysis, Domain Theory and Static Program
Analysis. He is now a freelance mathematician and computer scientist
working in Stirling.

Lahcen Ouarbya received a BSc in Physics from Caddi Ayad Univesity,
Marakech, Morocco and an Msc in Solid State Physics form St. Pet-
ersbourg Technical Univesity, St. Petersburg, Russia. He is currently
a PhD student under Sebastian Danicic at Goldsmiths College, Lon-
don University working on the Semantics of Program Slicing.

Martin Ward is a Visiting Research Fellow at De Montfort University
and Principal Consultant at Software Migrations Ltd. He works on the
theory and application of Program Transformations. He developed a
Wide Spectrum Language, called WSL, which is suitable for program
analysis and transformation.

1434
1435
1436
1437
1438
1439
1440
1441
1442
1443
1444
1445
1446
1447
1448
1449
1450
1451
1452
1453
1454
1455
1456
1457
1458
1459
1460
1461
1462
1463
1464
1465
1466
1467
1468