

Interprocedural Semantic Change-Impact Analysis using Equivalence Relations

Alex Gyori
University of Illinois
gyori@illinois.edu

Shuvendu K Lahiri
Microsoft Research
shuvendu@microsoft.com

Nimrod Partush
Technion
nimi@cs.technion.ac.il

Abstract—Change-impact analysis (CIA) is the task of determining the set of program elements impacted by a program change. Precise CIA has great potential to avoid expensive testing and code reviews for (parts of) changes that are refactorings (semantics-preserving). Existing CIA is imprecise because it is coarse-grained, deals with only few refactoring patterns, or is unaware of the change semantics.

We formalize the notion of change impact in terms of the trace semantics of two program versions. We show how to leverage equivalence relations to make dataflow-based CIA aware of the change semantics, thereby improving precision in the presence of semantics-preserving changes. We propose an *anytime* algorithm that allows applying costly equivalence relation inference incrementally to refine the set of impacted statements. We have implemented a prototype in SYMDIFF, and evaluated it on 322 real-world changes from open-source projects and benchmark programs used by prior research. The evaluation results show an average 35% improvement in the size of the set of impacted statements compared to standard dataflow-based techniques.

I. INTRODUCTION

Software constantly evolves to add and improve features, eliminate bugs, improve software design, etc. As software evolves faster than ever, it requires rigorous techniques to ensure that changes do not modify existing behavior in unintended ways. Several approaches have emerged to address this issue. Code review [37], regression testing [21], [47], test-suite augmentation [36], [42], [43], code contracts [5], [27], regression verification [22], [41] and verification modulo versions [35] are just a few of the approaches to ensure the quality of a change; they all directly benefit from change-impact analysis (CIA).

Change-Impact Analysis determines the set of program elements that may be impacted by a syntactic change. Traditional approaches are coarse-grained and operate at the level of types and classes [1], [2], or files [21] to retain soundness. Fine-grained techniques that aim to work at the level of statements are typically based on performing dataflow analysis [46] on one program to propagate the change along data and control flow edges [3], [10], [34]. Such techniques fail to take the *semantics of the change* into account; therefore, they cannot distinguish between changes that a user expects to have only local impact on existing code (e.g., a code refactoring) from ones that have substantial impact on existing code (e.g., changing the functionality or fixing a bug). The ability to distinguish changes whose impact is local (limited to the changed procedure or a few callers or callees within one or

two levels) can help with code review and regression-testing efforts. Changes with substantial impact can be prioritized for more rigorous code reviews and may require more testing.

In this paper, we aim to improve the precision of CIA by leveraging *equivalence relations* between variables of two programs across a change. At a high level, these equivalences help prune the flow of a change along the data or control flow edges of the changed program. To integrate such equivalences, we first formalize the notion of change impact precisely in terms of the trace semantics of two programs. Next, we show how to make CIA change-semantics aware by incorporating various equivalence relations into an interprocedural dataflow analysis. Since computing equivalence relations is expensive, we propose an *anytime algorithm* [49], [51] to incrementally compute equivalence relations.

A. Overview

Figure 1 shows a running example written in C. The example is inspired by real commits to Coreutils, in files `paste.c` [13] and `sort.c` [14]. The program has three changes. Two are semantics-preserving changes: (i) extracting the literal `'\n'` into the variable `line_delim` in the procedure `print_product_info` (lines 1, 4, 5) and (ii) replacing the conditional operator with a double negation in `locale_ok` (lines 22, 23)¹. The third change sets the `line_delim` variable to `'\0'` in the procedure `print_product_info` (lines 12, 13, 14), which impacts statements in `print_minor_version`. Assume for this example that all executions start from `print_product_info`. We claim that the only (syntactically unchanged) line that is *impacted by the changes* is the highlighted line 41 due to the semantic change at the callsite; a statement is impacted, intuitively, if the *sequence of values* it reads can differ when executing the two versions of the program in the same environment. For brevity, the example does not contain the definitions of the `setlocale` and `locale_format` procedures as well as the `LC_ALL` and `HEADER` constants as they are not impacted nor relevant. We will now analyze the change through the lens of a standard dataflow analysis [46] and traditional equivalence checking [22], [30] and then sketch our technique.

Dataflow: A dataflow analysis technique starts at the sources of change and propagates them through data and control

¹Negation in C coerces the values to 0 or 1.

```

1 + static unsigned char line_delim = '\n';
2 int print_product_info(int name, int version) {
3     int locale, printed = 0;
4 - print_header('\n');
5 + print_header(line_delim); // spurious arg impact
6     locale = locale_ok(); // spurious impact
7     if (name) {
8         printed = print_name(locale);
9     }
10    if (version && printed) {
11        printed = print_major_version(locale);
12 - printed = print_minor_version(locale, '\n');
13 + line_delim = '\0';
14 + printed = print_minor_version(locale, line_delim);
15    }
16    return printed;
17 }
18 void print_header(char delim) {
19     printf("%s%c", HEADER, delim);
20 }
21 int locale_ok() {
22 - return setlocale (LC_ALL, "") ? 1 : 0;
23 + return !!setlocale (LC_ALL, "");
24 }
25 int print_name(int locale) {
26     if (locale) {
27         printf("%s", locale_format("Coreutils"));
28         return 1;
29     }
30     return 0;
31 }
32 int print_major_version(int locale) {
33     if (locale) {
34         printf("%s", locale_format("8"));
35         return 1;
36     }
37     return 0;
38 }
39 int print_minor_version(int locale, char delim) {
40     if (locale) {
41         printf("%s%c", locale_format(".12"), delim);
42         return 1;
43     }
44     return 0;
45 }

```

Fig. 1. Example. The lines with $-$ and $+$ represent deleted and added lines, respectively.

edges (typically in the changed program). Dataflow techniques are not aware of the change semantics, and thus cannot exploit semantics-preserving changes. Initially, the call to `print_header` on line 5 has a change to its argument that marks all the statements in the procedure as impacted because they all depend on the changed argument. Next, the call to `locale_ok` on line 6 impacts the `locale` variable because of the change to the body of `locale_ok` and the data dependency of the return value on the change. This in turn will mark the input of `print_name` as impacted at line 8, which in turn flows to its output because the return value is control dependent on the input variable marked as impacted (a context-insensitive analysis will impact the return at *all* call sites to `print_name`). This impact through the return value will propagate to the call to `print_major_version` and `print_minor_version` because of the control dependency on `printed` and will impact all the statements in these procedures as well as all the returns at both call sites. Finally, the call to `print_minor_version` will impact all of the callee statements. A context-sensitive analysis does not help either because the body of `locale_ok` changes, which implies that the return value may change across the two versions. This is *imprecise* since the analysis is unable to determine that the statements in `print_name` and

`print_major_version` are not impacted.

Equivalence: A traditional interprocedural equivalence checking [22], [30] (that checks if two procedures have identical input-output behavior) will find that `locale_ok`, `print_name`, `print_header`, `print_major_version`, and `print_minor_version` have identical summaries. This is *unsound* for the question of impact analysis, as the statement of `print_minor_version` is impacted due to the change of print delimiter. This illustrates the difference between CIA and (traditional) equivalence checking: two procedures can be equivalent, but still impacted, because they may get called under different contexts and exhibit different behaviors.

Our approach: In this work, we present a *change-semantics aware* CIA that works as follows: it infers equivalence relations over variables and determines that the arguments at all call sites to `print_name` and `print_major_version` are equal across both versions and stops propagating impacts through their arguments. Further, `locale_ok` has an equivalent summary in the two versions (by using equivalence checking)—this ensures that two call sites with equal arguments return equal results. From these two facts, the technique infers (by simple dataflow) that arguments to `print_name` and `print_major_version` are not impacted and therefore the statements in `print_name` and `print_major_version` are not impacted. Thus, our approach precisely identifies the only unchanged impacted line as line 41.

B. Contributions

In this work, we make the following contributions:

- 1) We precisely formalize the set of statements *impacted* by a change, in terms of the trace semantics of two programs (§ III-A).
- 2) We make a dataflow-based CIA change-semantics aware by incorporating various equivalence relations (§ IV).
- 3) We describe an anytime algorithm that allows incrementally computing more equivalences to refine the analysis at the expense of time (§ IV-A).
- 4) We have implemented a prototype using SYMDIFF [30], [31], and evaluated our technique on 322 real-world changes collected from GitHub open-source projects and several standard benchmark programs used in prior research [26].

Verifiability: The prototype implementation of our tool is open source and available in the SymDiff repository <https://symdiff.codeplex.com/>. The usage options of the tool are described at <https://symdiff.codeplex.com/wikipage?title=Change-Impact>. Details about our experimental subjects are available at <http://mir.cs.illinois.edu/gyori/projects/cia.html>

II. BACKGROUND

For the rest of the paper, we will formalize the problem and our technique over a simple language. We can compile most features of general purpose imperative programming languages to our simple language; we discuss this in § II-B.

A. A Simple Language

A program consists of procedures represented as control-flow graphs, statements, and expressions.

Expressions: $e \in \text{Exprs}$ in the language are built up from constants, variables and operator applications:

$$e \in \text{Exprs} \quad :: \quad c \mid x \mid y \mid \dots \mid \text{op}(e_1, \dots, e_k)$$

Here c represents *constant* values of different types such as $\{\text{true}, \text{false}\}$ for Booleans, $\{\dots, -1, 0, 1, \dots\}$ for integers, and x denotes variables in scope. An operator **op** is a function or predicate symbol that can be uninterpreted or interpreted by some theories (e.g., $\{+, -, *, \leq, \geq, \dots\}$ by the theory of arithmetic). We represent a vector of variables and expressions using $\bar{x}$ and $\bar{e}$, respectively.

Statements: $st \in \text{Stmts}$ are comprised of *assign*, *assume*, *skip* and procedure *call* statements.

$$st \in \text{Stmts} \quad :: \quad x := e \mid \text{assume } e \mid \text{skip} \mid \text{call } x_1, x_2, \dots, x_k := f(e_1, e_2, \dots, e_m)$$

The argument to **assume** is a Boolean-valued expression, and a **skip** is a no-op. A call statement can have multiple return values and they are assigned to variables x_i at the call site.

Procedures: A procedure $f \in \text{Procs}$ is represented as a control-flow graph consisting of $(N_f, E_f, \text{In}_f, \text{Out}_f, \text{Vars}_f, n_f^e, n_f^x)$, where:

- N_f is a set of control-flow locations in f ,
- $E_f \subseteq N_f \times N_f$ is a set of edges over N_f denoting control-flow,
- In_f (respectively, Out_f) is the vector of *input* (respectively, *output*) formal of f . The output formal model return values and out parameters.
- Vars_f is the set of variables in the scope of f and includes In_f , Out_f , and local variables of f ,
- $n_f^e \in N_f$ (respectively, $n_f^x \in N_f$) is the unique *entry* (respectively, *exit*) node of f .

Let $N = \bigcup_{f \in \text{Procs}} N_f$ and $\text{Vars} = \bigcup_{f \in \text{Procs}} \text{Vars}_f$. Nodes and variables in a procedure f are often denoted by n_f and x_f respectively. For any node $n_f \in N_f$, we define the *readset* $R\text{Vars}(n_f)$ and *writeset* $W\text{Vars}(n_f)$ as the set of variables that are read and written to respectively in the statement at n_f .

A *program* $\text{Prog} \in \text{Programs}$ is a tuple $(\text{Procs}, \text{main}, \text{StmtAt})$ where (i) Procs is a set of procedures in the program, (ii) $\text{main} \in \text{Procs}$ is the entry procedure from which the program execution starts, and (iii) $\text{StmtAt} : N \rightarrow \text{Stmts}$ maps a node $n \in N$ in a procedure f to a *statement*. For any f , we assume that $\text{StmtAt}(n_f^x) = \text{skip}$.

B. Expressiveness

We can compile most constructs in general purpose imperative programming languages to our simple language. This follows the same principle as translators from languages such as C and Java to the Boogie language [4], [12], [20], [44].

Control flow: Loops can be automatically transformed into tail-recursive procedures [22], [30], [31]. We use $n_1 : st; \text{goto } n_2, n_3;$ to express that $\text{StmtAt}(n_1) = st$

and $\{(n_1, n_2)(n_1, n_3)\} \subseteq E_f$. A conditional statement **if** (e) st_1 **else** st_2 is modeled as:

$$n_1 : x := e; \text{goto } n_2, n_3; \\ n_2 : \text{assume } x; st_1; \text{goto } n_4; n_3 : \text{assume } \neg x; st_2; \text{goto } n_4;$$

where a fresh Boolean variable x captures the value of the condition e^2 . We assume that each node $n \in N_f$ has at most two successor nodes in E_f , where nodes with two successors correspond to conditional statements. The only use of an **assume** statement is to model a conditional statement. We refer to n_1 as a *branching* node with two successors in E with complementary expressions in **assume** statements.

Globals and heap: Richer data types such as arrays and maps can be modeled, e.g., array read $x[e]$ is modeled using $\text{sel}(x, e)$ and a write $x[e_1] := e_2$ is modeled using $x := \text{update}(x, e_1, e_2)$ [6]. Arrays are in turn used to model the heap in imperative programs and are standard in most software verification tools [12], [20], [44]. Additional internal non-determinism (e.g. read from file, network) is lifted as reads from immutable input arrays of *main*, making programs deterministic in our language [30]. We desugar the program's global variables (including the heap) as additional input and output formal arguments to a procedure. We transform each procedure into its *Static Single Assignment* (SSA) form [18], where a variable is assigned at exactly one program node.

C. Semantics

Let $\mathcal{V}$ denote the set of values that variables and expressions can evaluate to. Let $\theta \in \Theta$ be a *store* mapping variables to values in $\mathcal{V}$. For $x \in \text{Vars}$, we define $x \in \theta$ if x is a variable in the domain of θ . For $x \in \theta$, $\theta(x)$ denotes the value of variable x . The store $[x \rightarrow \nu]$ represents a singleton store that maps x to ν . The store $\theta|_{\text{Vars}_1}$ restricts the domain of the store to variables in Vars_1 . For stores θ_1 and θ_2 , the store $\theta_3 \doteq \theta_1 \oplus \theta_2$ is defined as follows for any variable $x \in \theta_1$ or $x \in \theta_2$:

$$\theta_3(x) = \begin{cases} \theta_2(x), & \text{if } x \in \theta_2 \\ \theta_1(x), & \text{otherwise} \end{cases}$$

The value of an expression $e \in \text{Exprs}$ ($\theta(e)$) is defined inductively on the structure of e (we omit it for brevity as it is fairly standard).

Calls: Let $cs \in (N \times \text{Vars}^* \times \Theta)^*$ be a *call stack* that is a sequence of tuples $\langle (n_0, \bar{r}_0, \theta_0), (n_1, \bar{r}_1, \theta_1), \dots \rangle$, where n_i is the i -th call site on the call stack (n_0 is the most recent), $\bar{r}_i$ and θ_i , respectively, are the vector of return actuals and the valuation of the local variables of the caller, at the corresponding call site. Let CS denote the set of all such call stacks, ϵ denotes an empty stack, and $(n, \bar{r}, \theta) :: cs$ denotes the *concatenation* operator.

Transition Relation: A *state* $\sigma \in \Sigma$ is a tuple $(n, \theta, cs) \in N \times \Theta \times \text{CS}$ that denotes a point in program execution where n is the current node being executed in a procedure f , θ is the valuation of variables in Vars_f and cs is the current call stack.

²The introduction of x simplifies determining if control flow is impacted by only inspecting the conditional node

A *state transition* denoted as $(n_f, \theta_1, cs_1) \rightsquigarrow (n_2, \theta_2, cs_2)$ is a relation over $\Sigma \times \Sigma$ holds only if:

- 1) $StmtAt(n_f) \doteq x := e$, $n_2 \in N_f$, $\theta_2 = \theta_1 \oplus [x \rightarrow \theta_1(e)]$, $(n_f, n_2) \in E_f$, and $cs_1 = cs_2$, or
- 2) $StmtAt(n_f) \doteq \mathbf{assume} \ e$, $n_2 \in N_f$, $\theta_1(e) = \mathbf{true}$, $(n_f, n_2) \in E_f$, $\theta_1 = \theta_2$ and $cs_1 = cs_2$, or
- 3) $StmtAt(n_f) \doteq \mathbf{skip}$, $n_f \neq n_f^x$, $n_2 \in N_f$, $(n_f, n_2) \in E_f$, $\theta_1 = \theta_2$ and $cs_1 = cs_2$, or
- 4) $StmtAt(n_f) \doteq \mathbf{call} \ \bar{r} := g(\bar{e})$. Let n be the unique successor of n_f in f , and $\bar{x}$ be the vector of input formals for g in $n_2 = n_g^e$, $cs_2 = (n, \bar{r}, \theta_1) :: cs_1$ and $\theta_2 = [\bar{x} \rightarrow \theta_1(\bar{e})]$, or
- 5) $StmtAt(n_f) \doteq \mathbf{skip}$, $n_f = n_f^x$, $cs_1 \doteq (n_g, \bar{r}, \theta_3) :: cs_3$. Let $\bar{y}$ be the vector of output formals for f in $n_2 = n_g$, $\theta_2 = (\theta_3 \oplus [\bar{r} \rightarrow \theta_1(\bar{y})])|_{Vars_g}$, $cs_2 = cs_3$.

A transitive edge $\sigma_0 \rightsquigarrow^* \sigma_n$ exists if $\sigma_n \equiv \sigma_0$ or there exists a sequence of transitions $\sigma_0 \rightsquigarrow \sigma_1, \dots, \sigma_{n-1} \rightsquigarrow \sigma_n$, where $\sigma_i \rightsquigarrow \sigma_{i+1}$, for all $i \in [0, \dots, n)$. For a procedure f , we denote the input-output *transition relation* $\Omega_f \doteq \{(\theta_1, \theta_2) \mid (n_f^e, \theta_1, \epsilon) \rightsquigarrow^* (n_f^x, \theta_2, \epsilon)\}$.

Execution Traces: An *execution trace* τ is a (possibly infinite) sequence of states $\langle \sigma_0, \sigma_1, \dots \rangle$, where $\sigma_i \rightsquigarrow \sigma_{i+1}$, for any adjacent pair of states in the sequence. For a trace τ and a node $n \in N$, $\tau|_n$ denotes the (maximal) subsequence of τ containing states of the form $(n, _, _)$. For such a trace τ of length at least $i+1$, $\tau[i]$ denotes the state at position i (namely σ_i). For any procedure f , let Γ_f be the set of *maximal* traces of f . That is, Γ_f is the set of all traces τ such that (i) $\tau[0] \doteq (n_f^e, _, \epsilon)$, and (ii) either (a) the final state σ_n has no successors, or (b) the trace is non-terminating. Traces with no successors can either terminate *normally* in a state $(n_f^x, _, \epsilon)$, or could be *blocked* due to no successors in E or due to an unsatisfied **assume** statement. For a store $\theta \in \Theta$, we denote $\tau_f(\theta)$ as the maximal trace (due to determinism) of f that starts in a store θ .

III. PROBLEM STATEMENT

In this section, we formalize the problem of *semantic change-impact analysis* and provide a simple solution based on dataflow-based static analysis.

A. Representing Changes

We denote $Prog^1, Prog^2 \in Programs$ as two versions of a program. Similarly $\sigma^i, \theta^i, \tau^i, Procs^i, main^i, StmtAt^i$ denote entities for $Prog^i$, without making $Prog^i$ explicit.

To simplify the formulation we assume the two programs in a *normalized* form, where (i) each procedure in $Procs^1$ has a corresponding procedure in $Procs^2$ and vice versa, and (ii) for each $f \in Procs^i$, the vector of variables in $Vars_f$, and the set of nodes N_f (but not necessarily E_f) are identical with the ones in the corresponding procedure. We can easily preprocess the programs to obtain their normalized form, by introducing additional procedures, variables (uninitialized) and nodes. Finally, for any missing node n , we add an unreachable node in N_f with a **skip** statement and empty successor list.

Diffing: Given the two versions, a diffing algorithm produces a mapping between nodes in the two programs. We assume we

are given a sound diff algorithm to label the sources of change. A diff algorithm is sound if it produces a partial function $\pi : N_1 \rightarrow N_2$ such that:

- 1) π is a partial bijection³ and $StmtAt(n_f) = StmtAt(\pi(n_f))$.
- 2) For any two traces $\tau^1 \doteq \tau_{main}^1(\theta)$ in $Prog^1$ and $\tau^2 \doteq \tau_{main}^2(\theta)$ in $Prog^2$, τ^1 only executes statements in $Dom(\pi)$ iff τ^2 only executes statements in $Im(\pi)$.
- 3) For any two traces $\tau^1 \doteq \tau_{main}^1(\theta)$ in $Prog^1$ and $\tau^2 \doteq \tau_{main}^2(\theta)$ in $Prog^2$, where τ^1 only executes statements in $Dom(\pi)$ or τ^2 only executes statements in $Im(\pi)$, then $\tau^1 = \tau^2$.

The mapped nodes $MAPPED \doteq Dom(\pi) \cup Im(\pi)$ underapproximate the set of nodes that are syntactically unchanged. Intuitively, if a program executes only statements in MAPPED then the program behaves the same in both versions; statements that are not in MAPPED are the sources of change.

We describe for illustrative purposes a simple diffing algorithm which is sound. The algorithm proceeds to produce a mapping π as follows: Let $Procs^\Delta \subseteq Procs$ be the set of procedures that have some syntactic change. Any node not in $f \in Procs^\Delta$ is trivially mapped as the control-flow graphs are identical in the two versions. Any node in $f \in Procs^\Delta$ is conservatively treated as not mapped. Our formulation is parameterized by a diff algorithm which can either be based on text [50] or more sophisticated notions such as abstract syntax trees [17] or program-dependency-graphs [29] as long as they satisfy the soundness criteria.

B. Semantic Change Impact

We can now state the meaning of a node being impacted by a program change, in terms of the trace semantics of the two programs and the set MAPPED.

For a sequence of states $\bar{\sigma}$ and a variable $x \in Vars$, $\bar{\sigma}|_x \in (V \cup \{\perp\})^*$ denotes the sequence of values $\bar{v}$ with the same length as $\bar{\sigma}$, and

$$\nu_i = \begin{cases} \theta(x), & \sigma_i \doteq (_, \theta, _) \text{ and } x \in \theta \\ \perp & \text{otherwise} \end{cases}$$

Definition 1 (Impacted nodes): Given $Prog^1, Prog^2$ and MAPPED, a node $n \in N^1 \cup N^2$ is *impacted* if either $Impacted(n, Prog^1, Prog^2, \pi)$ or $Impacted(\pi(n), Prog^2, Prog^1, \pi^{-1})$ holds, where π^{-1} is the inverse. N^i is the corresponding N for $Prog^i$.

We define $Impacted(k, Prog^a, Prog^b, \varpi)$:

- 1) $k \notin Dom(\varpi)$, or
- 2) there exists a store θ , pair of traces $\tau^a \doteq \tau_{main}^a(\theta)$ for $Prog^a$ and $\tau^b \doteq \tau_{main}^b(\theta)$ for $Prog^b$, and a variable $x \in RVars(n)$ such that $(\tau^a|_x)|_k \neq (\tau^b|_{\varpi(k)})|_x$.

We conservatively treat any unmapped node as impacted. A mapped node n is not impacted if the sequence of values of variables in $RVars(n)$ is identical for any two execution traces τ^a (in $Prog^a$) and τ^b (in $Prog^b$) starting from a common input store θ to *main*. Note that for our low-level language, the $RVars(n)$ of a statement often includes

³A partial bijection is a partial function that is injective when defined and (trivially) surjective when restricted to its image [23].

Predicate name	Definition
BRANCHINGNODE(n)	if n is a branching node
CONTROLDEPENDENT(n_2, n_1)	if n_2 is control-dependent on n_1 [19]
CALLSITE(n, f, g)	if $SimAt(n)$ is a call to f within a caller g .
INFORMAL(x, i, f)	if x is the i -th input formal of f
OUTFORMAL(x, i, f)	if x is the i -th output formal of f
INACTUAL(e, i, f, n)	if the expression e is the i -th actual argument to a call to f at a callsite n
OUTACTUAL(r, i, f, n)	if the variable r receives the i -th output formal to a call to f at a callsite n

TABLE I
PREDICATES USED FOR DATAFLOW ANALYSIS.

the state of the heap and address being written to. For example, the C# statement $x.length = y$ gets translated to $n : Length := update(Length, x, y)$, (where $Length$ is an array variable representing the state of $length$ field/attribute in all objects) with $RVars(n) = \{Length, x, y\}$.

C. Dataflow-based Change-Impact Analysis

In this section, we describe *Dataflow-based Change-Impact Analysis* (DCIA), a *change semantics unaware* static analysis that provides a conservative estimate of the set of impacted nodes. The static analysis is an interprocedural dataflow analysis [46] that starts with a program $Prog^i$ ($i \in 1, 2$) and a conservative estimate of the syntactically-changed nodes, nodes not in MAPPED, and returns an upper bound on the set of (a) impacted nodes, (b) impacted variables, and (c) output variables whose summary may have changed.

Predicates: Table I defines some straightforward predicates used in the inference rules. The $OUTACTUAL(r, i, f, n)$ predicate holds when the i^{th} return value is assigned to variable r , at the call to f from the node n (note that we allow multiple return values); we call r the output actual to differentiate it from the i^{th} output formal inside the callee. For $CONTROLDEPENDENT(n_2, n_1)$, a node n_2 is control-dependent on node n_1 iff (i) there exists a path from n_1 to n_2 s.t. every node in the path other than n_1 and n_2 is *post-dominated* by n_2 , and (ii) n_1 is not post-dominated by n_2 [19].

Dependency: Figure 2 describes a set of inference rules to compute two relations $DEPENDSONVAR$ and $DEPENDSONNODE$. For a pair of variables $x, y \in Vars_f$ such that y is either data- or control-dependent on x , then $DEPENDSONVAR(y, x, f)$ holds. Similarly, a node $n \in N_f$ and a variable y such that y is data or control dependent on a variable x that is updated at n , $DEPENDSONNODE(y, n, f)$ holds. An inference rule (e.g. $DEPENDS-NODE$) lists a set of antecedents (above the line) and the consequent (below the line). Applying an inference rule results in adding a tuple to the relation in the consequent (e.g. $DEPENDSONNODE$). The inference rules are applied repeatedly until a fix-point is reached.

Most of the inference rules are straightforward encoding of program data- and control flow. The rule $CONTROL-DEPENDS$ expresses that if n_1 is a branching node, whose condition depends on x and y is written in a control-dependent node n_2 , then y depends on x . The rule $SUMMARY-DEPENDS$ captures the dependency of an actual return r

DEPENDS-ENTRY
$x \in In_f$
$DEPENDSONVAR(x, x, f)$
DEPENDS-WRITE
$x \in RVars(n) \quad y \in WVars(n) \quad n \in N_f$
$DEPENDSONVAR(y, x, f)$
DEPENDS-TRANSITIVE
$DEPENDSONVAR(y, x, f) \quad DEPENDSONVAR(x, z, f)$
$DEPENDSONVAR(y, z, f)$
CONTROL-DEPENDS
$BRANCHINGNODE(n_1) \quad x \in RVars(n_1)$
$CONTROLDEPENDENT(n_2, n_1) \quad y \in WVars(n_2)$
$DEPENDSONVAR(y, x, f)$
SUMMARY-DEPENDS
$CALLSITE(n, f, g)$
$OUTACTUAL(r, i, f, n) \quad OUTFORMAL(y, i, f)$
$INFORMAL(x, j, f) \quad DEPENDSONVAR(y, x, f)$
$INACTUAL(e, j, f, n) \quad w \in RVars(e)$
$DEPENDSONVAR(r, w, g)$
DEPENDS-NODE
$x \in WVars(n) \quad n \in N_f \quad DEPENDSONVAR(y, x, f)$
$DEPENDSONNODE(y, n, f)$

Fig. 2. Inference rules for computing $DEPENDSONVAR$ and $DEPENDSONNODE$. The input is a program $Prog$.

on a variable w passed as an argument to f in a caller g , where w indirectly flows to r through a procedure call to f . For this callsite, the i -th output formal y (which is assigned to the output actual r) is dependent on the j -th input formal x , which in turn is assigned the actual e at the callsite.

Impact Analysis: Figure 3 describes a set of inference rules to compute the set of nodes that are impacted in either program. For this section, we will ignore the **highlighted** antecedents (they become relevant in § IV where we describe how we incorporate change semantics). The rules take as input a program (either $Prog^1$ or $Prog^2$), the set of mapped nodes MAPPED and precomputed relations $DEPENDSONNODE$ and $DEPENDSONVAR$ for the particular program. They produce the relations IMPACTEDNODE, IMPACTEDVAR and IMPACTEDSUMM that represent an upper bound on the set of impacted nodes, impacted variables and impacted variable summaries, respectively. Most rules are self-explanatory; we describe the main rules relevant to the interprocedural reasoning. Note that we do not have a special rule for control-flow impact, since $DEPENDSONVAR$ already captures control flow dependency.

For an output formal $y \in Out_f$, the summary (input-output dependency) may change either when (i) y depends on a variable updated at an unmapped node $n \in N_f$ (expressed by $IMPACT-SUMMARY$), or (ii) y depends on a variable w that stores the output formal x of g at a callsite in node n , and the summary of x has changed in g (expressed by $IMPACT-SUMMARY-PROP$).

The $CALL-IMPACT$ rule says that an input formal x in f can be impacted if the corresponding actual argument e

$\frac{\text{SYNT-CHANGED} \quad n \notin \text{MAPPED}}{\text{IMPACTEDNODE}(n)}$	$\frac{\text{NODE-2-VAR} \quad \text{IMPACTEDNODE}(n) \quad x \in WVars(n)}{\text{IMPACTEDVAR}(x)}$	$\frac{\text{VAR-2-EXPR} \quad \text{IMPACTEDVAR}(x) \quad x \in RVars(e)}{\text{IMPACTEDEXPR}(e)}$
$\frac{\text{VAR-2-NODE} \quad \text{IMPACTEDVAR}(x) \quad x \in RVars(n)}{\text{IMPACTEDNODE}(n)}$		
$\frac{\text{IMPACT-SUMMARY} \quad \text{OUTFORMAL}(y, i, f) \quad \text{DEPENDSONNODE}(y, n, f) \quad n \notin \text{MAPPED} \quad \neg \text{SUMMARYEQUIV}(y, f)}{\text{IMPACTEDSUMM}(y, f)}$		
$\frac{\text{IMPACT-SUMMARY-PROP} \quad \text{OUTFORMAL}(x, j, g) \quad \text{IMPACTEDSUMM}(x, g) \quad \text{OUTFORMAL}(y, i, f) \quad \text{CALLSITE}(n, g, f) \quad \text{OUTACTUAL}(w, j, g, n) \quad \text{DEPENDSONVAR}(y, w, f) \quad \neg \text{SUMMARYEQUIV}(y, f)}{\text{IMPACTEDSUMM}(y, f)}$		
$\frac{\text{CALL-IMPACT} \quad \text{CALLSITE}(n, f, g) \quad \text{INACTUAL}(e, i, f, n) \quad \text{IMPACTEDEXPR}(e) \quad \text{INFORMAL}(x, i, f) \quad \neg \text{PREEQUIV}(x, f)}{\text{IMPACTEDVAR}(x)}$		
$\frac{\text{RETURN-IMPACT} \quad \text{CALLSITE}(n, f, g) \quad \text{OUTACTUAL}(r, i, f, n) \quad \text{OUTFORMAL}(y, i, f) \quad \text{IMPACTEDSUMM}(y, f)}{\text{IMPACTEDVAR}(r)}$		
$\frac{\text{SUMMARY-IMPACT} \quad \text{CALLSITE}(n, f, g) \quad \text{OUTACTUAL}(r, i, f, n) \quad \text{OUTFORMAL}(y, i, f) \quad \text{INFORMAL}(x, j, f) \quad \text{DEPENDSONVAR}(y, x, f) \quad \text{INACTUAL}(e, j, f, n) \quad \text{IMPACTEDEXPR}(e) \quad \neg (\text{PREEQUIV}(x, f) \wedge \text{SUMMARYEQUIV}(y, f))}{\text{IMPACTEDVAR}(r)}$		

Fig. 3. Inference rules for dataflow based change-impact analysis. The **highlighted** antecedents are relevant for change-semantics aware analysis.

at a callsite is impacted. RETURN-IMPACT considers the case when the variable summary for the corresponding output formal y is impacted. SUMMARY-IMPACT considers the case when the actual argument expression e (passed for the formal x of f) is impacted, and y depends on the value of x in f . Our analysis preserves context-sensitivity as it does not impact a return value simply because the corresponding output formal is impacted in some context.

The algorithm DCIA does the following:

- 1) Takes as input $Prog^1, Prog^2$ and MAPPED .
- 2) Applies the inference rules in Figure 3 on $Prog^i$ to generate IMPACTEDNODE^i , IMPACTEDVAR^i , IMPACTEDSUMM^i until a fix-point is reached.
- 3) Returns the tuple $(\bigcup_i \text{IMPACTEDNODE}^i, \bigcup_i \text{IMPACTEDVAR}^i, \bigcup_i \text{IMPACTEDSUMM}^i)$.

The following theorem states the soundness of the dataflow analysis DCIA.

Theorem 1 (Soundness): Given $Prog^1, Prog^2 \in \text{Programs}$ and $\text{MAPPED} \subseteq N$, (a) DCIA terminates, and (b) for any $n \notin \text{IMPACTEDNODE}$, n is not an impacted node with respect to MAPPED (according to Definition 1).

Consider the changes in Figure 1 at line 22; the procedure `locale_ok` has an impacted summary because its return variable depends on a node that is syntactically changed, i.e., is not in MAPPED . This causes the line 6 and the variable `locale` to be marked as impacted because of the rule IMPACT-SUMMARY. Impacts are propagated interprocedurally by the rule CALL-IMPACT to all calls that take `locale`

as an argument, i.e., `print_name`, `print_major_version`, and `print_minor_version`. Similarly, using the same rule, the body of `print_header` is impacted by the changed argument ‘\n’ changed to the variable `line_delim` on line 4. The propagation through calls further impacts their entire body because of the data and control dependency on the impacted argument (by the rules NODE-2-VAR and VAR-2-NODE which propagate impact through both control- and data-dependency relying on the predicate `DEPENDSONVAR`).

IV. INCORPORATING CHANGE SEMANTICS

In this section, we make the DCIA algorithm *change-semantics aware*. In other words, the analysis takes into account also the exact semantics of the change, in addition to the set of nodes MAPPED that may have been syntactically changed. We inject the change-semantics by leveraging equivalence relationships between variables and procedure summaries in the two programs $Prog^1$ and $Prog^2$.

Let us define the following semantic equivalences for a variable over $Prog^1$ and $Prog^2$.

Definition 2 (PREEQUIV): $\text{PREEQUIV}(x, f)$ holds for an input formal $x \in \text{In}_f$ if for all stores θ , and for every pair of traces $\tau^1 \doteq \tau_{\text{main}}^1(\theta)$ and $\tau^2 \doteq \tau_{\text{main}}^2(\theta)$, $(\tau^1|_{n_f^e}) \downarrow_x = (\tau^2|_{n_f^e}) \downarrow_x$.

Intuitively, $\text{PREEQUIV}(x, f)$ holds for an input formal x of f if any two executions starting from the same input θ to *main* call f with the same sequence of values of x . For example, in Figure 1 the equivalences $\text{PREEQUIV}(\text{delim}, \text{print_header})$, $\text{PREEQUIV}(\text{locale}, \text{print_name})$,

$\text{PREEQUIV}(\text{locale}, \text{print_major_version})$, and $\text{PREEQUIV}(\text{locale}, \text{print_minor_version})$ hold. In contrast, $\text{PREEQUIV}(\text{delim}, \text{print_minor_version})$ does not hold, because of different values for *delim* ‘\n’ and ‘\0’ respectively, at the call-site in *print_product_info*.

Let us define $\text{Deps}(y)$ as the set of variables x such that $\text{DEPENDSONVAR}(y, x, f)$ in either Prog^1 or Prog^2 . For two stores θ_1 and θ_2 defined over same set of variables, we denote $\theta_1 =_{\text{vars}_1} \theta_2$ to mean $\theta_1(x) = \theta_2(x)$ for every $x \in \text{Vars}_1$.

Definition 3 (SUMMARYEQUIV). $\text{SUMMARYEQUIV}(y, f)$ holds for an output formal $y \in \text{Out}_f$ if $(\theta_1, \theta_2) \in \Omega_f$ in Prog^i and $\theta_1 =_{\text{Deps}(y)} \theta_3$, then $(\theta_3, \theta_4) \in \Omega_f$ is in Prog^j ($j \neq i$) and $\theta_2(y) = \theta_4(y)$.

Intuitively, if the versions of f are executed from stores θ_1 and θ_3 where $\theta_1 =_{\text{Deps}(y)} \theta_3$, then either both procedures do not terminate, or the value of y after executing f is identical on exit. In our example in Figure 1, all procedures are equivalent except *print_product_info*, i.e., in this case $\text{SUMMARYEQUIV}(\text{line_delim}, \text{print_product_info})$ does not hold since in one version the value of *line_delim* at the end of the execution is “\0” while in the other it is undefined.

Figure 3 with the highlighted parts provides a refinement to the dataflow analysis to incorporate change semantics. In addition to the MAPPED, the algorithm now takes as input pre-computed relations PREEQUIV and SUMMARYEQUIV. In this section, we assume an oracle that provides these relations; we provide one implementation later (§ V-A). The highlighted facts strengthen the antecedent of a rule and prevent it from being applicable in some contexts. For example, the strengthened CALL-IMPACT prevents an input formal x from being impacted if $\text{PREEQUIV}(x, f)$ holds. Similarly, the strengthened IMPACT-SUMMARY prevents a summary for y from impact if we know that $\text{SUMMARYEQUIV}(y, f)$ holds. The strengthened SUMMARY-IMPACT is now applicable only when either (i) the formal x does not satisfy PREEQUIV or (ii) the summary for y does not satisfy SUMMARYEQUIV.

We denote the new change-semantics aware algorithm as *Semantic Dataflow-based Changed Impact Analysis (SEM-DCIA)*.

Theorem 2 (Soundness). Given $\text{Prog}^1, \text{Prog}^2 \in \text{Programs}$, MAPPED, PREEQUIV, and SUMMARYEQUIV, (i) SEM-DCIA terminates, and (ii) for any $n \notin \text{IMPACTEDNODE}$, n is not an impacted node with respect to MAPPED (from Definition 1).

A. Anytime Algorithm

The SEM-DCIA algorithm assumes an oracle to compute the PREEQUIV and SUMMARYEQUIV relations. Computing such equalities typically require constructing the product of the two programs Prog^1 and Prog^2 and inferring equivalence relations over the product program [31]. Such inference algorithms typically have high complexity and therefore it is wise to apply them prudently. In this section, we make a simple observation that allows us to interleave SEM-DCIA and inference of PREEQUIV and SUMMARYEQUIV in a single framework.

To exploit the change semantics, it is often useful to apply equivalence relation inference only in the vicinity of actual syntactic changes.

```

void main(int x)
{
- f1(x);
+ f1(x+0);
}
void f1(int x)
{
  f2(x+1);
}

void f2(int x)
{
  f3(x+2);
}
...
void fn(int x)
{
  ;
}

```

Fig. 4. Motivating example for anytime algorithm.

Consider the example in Figure 4 to make the intuition clear. Applying DCIA will result in impacting all the nodes in the program as follows. The modified call node for f_1 in *main* is not in MAPPED, which impacts input formal x of f_1 . This in turn impacts the call to f_2 and so on. On the other hand, we can observe that PREEQUIV and SUMMARYEQUIV hold for each of the procedures because the change does not propagate outside the changed statement.

For Figure 4 it suffices to infer the equivalences on *main* while abstracting the rest of the procedures from the expensive equivalence analysis. Considering f_1 has all callsites inside *main* and that it does not have an impacted summary by rule IMPACT-SUMMARY after DCIA suffices to determine that $\text{PREEQUIV}(x, f_1)$ holds. This information can be fed to SEM-DCIA which will prune the impact for the input parameter of f_1 which will prune the remaining impacts when performing a pure dataflow analysis. Thus, we obtain a precise change-impact analysis by applying the equivalence inference only on a small subset of the procedures in the program. Similarly, in Figure 1 it suffices to analyze only the syntactically changed procedures and abstract away the others to obtain the most precise result; this is not the case in general because to infer the PREEQUIV we need all call sites to be in scope, not only the syntactically changed procedures.

Algorithm 1 (SEM-DCIA-ANYTIME) provides an *anytime* algorithm that performs the integration. The algorithm takes as an additional input Procs^Δ , the set of syntactically changed procedures. It outputs a set of nodes *impNds* that overapproximates the set of impacted nodes. We term the algorithm anytime [16], [49], [51] because the algorithm can be stopped at any time after the first call to SEM-DCIA to obtain a conservative bound for the impacted nodes.

The algorithm starts with invoking SEM-DCIA on the two programs with an empty set of equivalences in *EQ* (line 4); this is identical to calling DCIA. The return values provide a conservative measure on impacted variables, nodes and summaries respectively (Theorem 1). The algorithm implements a loop (line 6) where it increases the frontier of procedures Procs' around Procs^Δ that are analyzed for inferring equivalences in *InferEquivs* (line 13). Lines 7 and 8 construct equivalences from the provably non-impacted variables and summaries. These equivalences are added to *EQ* in line 9. *ProcsWithin* returns all procedures that can reach or be reached from Procs^Δ within a call stack of depth k ; k is incremented with each iteration of the loop. *AbstractProcs* abstracts all procedures outside Procs' ; it only retains the knowledge of whether any procedure $f \in \text{Procs}'$ has additional call sites outside Procs' - this determines whether PREEQUIV can be

Algorithm 1: SEM-DCIA-ANYTIME

Input: $Prog^1, Prog^2 \in Programs$
Input: $Procs^\Delta \subseteq Procs$
Input: $MAPPED \subseteq N$
Output: $impNds \subseteq N$

```

1 begin
2    $k \leftarrow 0$ ;
3    $EQ \leftarrow (\emptyset, \emptyset)$ ;
4    $(impNds, impVars, impSumms) \leftarrow$ 
      SEM-DCIA( $Prog^1, Prog^2, MAPPED, EQ$ );
5    $Procs' \leftarrow Procs^\Delta$ ;
6   while  $Procs' \subset Procs$  do
7      $prEQ \leftarrow \{(x, f) \mid x \in In_f \text{ and } x \notin impVars\}$ ;
8      $smEQ \leftarrow \{(x, f) \mid x \in Out_f \text{ and } (x, f) \notin impSumms\}$ ;
9      $EQ \leftarrow EQ + (prEQ, smEQ)$ ;
10     $Procs' \leftarrow$ 
      ProcsWithin( $Procs^\Delta, Prog^1, Prog^2, k$ );
11     $Prog_k^1 \leftarrow AbstractProcs(Prog^1, Procs \setminus Procs')$ ;
12     $Prog_k^2 \leftarrow AbstractProcs(Prog^2, Procs \setminus Procs')$ ;
13     $EQ \leftarrow InferEquivs(Prog_k^1, Prog_k^2, EQ)$ ;
14     $(impNds, impVars, impSumms) \leftarrow$ 
      SEM-DCIA( $Prog^1, Prog^2, MAPPED, EQ$ );
15     $k++$ ;
16 return  $impNds$ ;
```

inferred for a procedure. `InferEquivs` is invoked with a set of equivalences in EQ on the smaller programs $Prog_k^i$. The final call to SEM-DCIA is used to compute the more refined set of impacted variables, nodes and summaries based on the equivalences discovered from `InferEquivs`. The loop terminates when $Procs'$ consists of the entire program; at this point `InferEquivs` has already looked at the entire program and no new equivalences will be discovered in line 13.

Let us denote $SEM-DCIA_k$ as an instantiation of the algorithm SEM-DCIA-ANYTIME that terminated after the loop is executed exactly $k + 1$ times. We also denote $SEM-DCIA_\infty$ if the loop terminates normally after $Procs'$ equals $Procs$.

Theorem 3 (Soundness): Given $Prog^1, Prog^2 \in Programs$, $MAPPED$, and $Procs^\Delta$, if $SEM-DCIA_k$ terminates then for any $n \notin impNds$, n is not an impacted node with respect to $MAPPED$ (according to Definition 1).

V. IMPLEMENTATION AND EVALUATION

A. Implementation

We implemented our SEM-DCIA analysis for C programs, but our analysis works on the intermediate verification language Boogie [4]. We leverage SMACK [44] to convert LLVM bytecode to Boogie programs.

Diffing: For our initial implementation, we leveraged `diff` over C files to produce the source of changes, i.e., nodes not in $MAPPED$. However, `diff` does not satisfy the soundness criteria for `diff` (Section III-A) because of changes in macros, data

Project Name	# Version Pairs	SLOC		LOC Changed	
		min	max	min	max
flingfd	2	142	146	2	14
histo	8	617	624	1	6
mdp	91	135	1616	1	402
theft	2	1672	1838	2	328
tinyvm	61	425	903	1	328
print_tokens	5	478	480	1	8
print_tokens2	10	397	402	1	6
replace	32	509	516	1	15
schedule	9	290	294	2	4
space	38	6180	6205	1	42
tcas	41	136	140	2	16
tot_info	23	346	347	2	3

TABLE II
SUMMARY OF PROJECTS USED AS EVALUATION SUBJECTS

structures, control-flow changes, etc.; we therefore conservatively consider all nodes in a changed procedure as sources of impacts. Although this can overapproximate the initial source of impact, the use of equivalences in SEM-DCIA allows us to prune the spurious impacts from escaping the syntactically-changed procedures; All our code and scripts are available in the SYMDIFF repository at: <https://syndiff.codeplex.com/>.

Inference: We used SYMDIFF to construct a product program and infer valid PREEQUIV and SUMMARYEQUIV. Given $Prog^1$ and $Prog^2$, SYMDIFF generates a product program $Prog^{1 \times 2}$ that defines a procedure $f^{1 \times 2}$ for every f and $\pi(f) \in Procs^i$. For the product program $Prog^{1 \times 2}$, one can leverage any of the (single program) invariant generation techniques to infer preconditions, postconditions (including two-state postconditions) on $f^{1 \times 2}$. Such invariants are *relational* in that they are over the state of two programs $Prog^1$ and $Prog^2$, and include equivalences relations such as PREEQUIV (preconditions of $f^{1 \times 2}$) and SUMMARYEQUIV (summary of $f^{1 \times 2}$). To ensure our inferred equivalences are valid we require the programs to be equi-terminating [25]; this is an area of future work – for now we assume that changes do not introduce non-termination. We modified SYMDIFF to add candidates for inferring summaries and take as input cheaply-inferred equalities from DCIA. More details can be found in our extended report [24].

B. Evaluation

In this section we demonstrate the effectiveness of our approach on GitHub projects with real program changes and also standard benchmark programs with artificial changes. We show that our semantic based analysis, SEM-DCIA improves on DCIA by reducing the size of the impacted set. The size of the impacted set is a proxy metric for the effort necessary to perform many software engineering tasks such as code review and testing.

We analyze 164 actual changes consisting of refactorings, feature additions, buggy changes, and bug fixes from 5 projects from the GitHub repository. The projects, number of versions used, their size in source lines of code (SLOC), and corresponding change sizes (in number of C source lines changed) are summarized in Table II. Our subjects are applications written in C such as a virtual machine program (tinyvm), a histogram creator (histo), a markdown presentation tool

(mdp), a file-descriptor management library (flingfd) and a test-generation library (theft). In addition, we also include 6 standard benchmark programs widely used by prior research on regression testing [26]. These benchmarks consist of 158 manually introduced changes representing non-trivial and hard to detect bugs. Our projects are sized between 142 lines of source code and 6205 (SLOC). The changes in our projects vary in size between very small changes, consisting of single line changes and larger ones, consisting of over 400 lines (most of our changes are on the small end of this spectrum).

For our experiments, we first compare SEM-DCIA against DCIA to study the impact of adding change-semantics to the impact analysis (§ V-C). Next, we evaluate the cost-precision tradeoff of the anytime algorithm SEM-DCIA-ANYTIME (§ V-D). Finally, we present several representative examples discovered while applying our tool (§ V-E).

C. Change-Semantic Aware Analysis

Table III shows the results of running our SEM-DCIA analysis on each of our subjects. For each change, we measure the number of lines reported as impacted by dataflow analysis (columns DCIA Impact) and also by SEM-DCIA (columns SEM-DCIA_∞). The columns SEM-DCIA_i denote various bounds for SEM-DCIA-ANYTIME and results will be discussed in § V-D. We report for each project the minimum and maximum number of impacted lines (min, max), and for the SEM-DCIA analysis we report also the average reduction of the size of the impacted set. Note that SEM-DCIA analysis always reports a subset of the set reported by the non-semantic analysis. We also report the average analysis time in seconds for the non-semantic analysis and for the SEM-DCIA analysis.

Our evaluation shows that on average, the change-aware analysis reduces the size of the impacted set by 35%. The overhead of performing full semantic analysis on the entire program is on median 19x, ranging between 3x and 67x. While the semantic analysis results at ∞ level represent the most precise analysis our technique achieves, it is quite expensive, and it even times-out for our largest program (e.g. *space*). For example in the theft project the reduction achieved by SEM-DCIA_∞ is 77% but with a 64x overhead. This motivates the need for an incremental analysis, whose results can be obtained faster.

Imprecision: Our manual inspection of results reveals three broad classes for nodes classified as impacted: (i) nodes in syntactically changed procedures, (ii) SYMDIFF’s inability to match loops as it relies on syntactic position in AST (this can be fixed by better matching heuristics), (iii) SMACK represents all aliased addresses accessing a field using a single map; writing to one location destroys equivalences on the map variables (need more refined conditional equivalences [28]).

D. Incremental Analysis

Table III shows the results of varying the bound on k for the SEM-DCIA-ANYTIME. The first iteration SEM-DCIA₀ corresponds to semantically analyzing only the syntactically-changed procedures; the second iteration SEM-DCIA₁ corre-

Analysis	Min	Max	Reduction	Time
DCIA	20	2851	<i>n.a.</i>	59.45
SEM-DCIA ₀	14	2816	31.87%	798.96
SEM-DCIA ₁	14	2816	36.71%	895.14
SEM-DCIA ₂	14	2816	40.56%	1300.43
SEM-DCIA ₃	14	2808	43.96%	1900.03
SEM-DCIA _∞	<i>n.a.</i>	<i>n.a.</i>	<i>n.a.</i>	<i>timeout</i>

TABLE IV
ANALYSIS RESULTS FOR SPACE

```
void draw_histogram(int data[], int len) {
    ...
    + int xbarw = 5;
    ...
    while (y-- > 0) {
        - term_move_to(x * 5 + xpad + 3,
        + term_move_to(x * xbarw + xpad + 3,
            y - 1 + h + ypad);
        ...
    }
}
```

Fig. 5. Change illustrating an extract constant to variable in histo commit c723a4

sponds to analyzing the procedures at distance at most one from the syntactically changed procedures. The results show that even SEM-DCIA₀ provides benefits, pruning the impacted set by 22% on average. The overhead is reduced compared to the full analysis (9x). Similarly, results for SEM-DCIA₁ show the analysis is effective. For the case of theft the improvement is from 61% (SEM-DCIA₀) to 77% (SEM-DCIA_∞), at the cost of overhead increase from 8x to 64x.

In general, we foresee that the anytime analysis is most useful for cases where it is prohibitive to run the full algorithm because of time constraints. This is best illustrated for the case of *space* (we used a timeout of one hour). Table IV shows the first four levels for *space* (two more iteration beyond the ones in Table III); performing the analysis incrementally is still valuable even upto $k = 3$; the first iteration already provides big benefits on top of the non-semantic analysis, while the following iterations display a smooth improvement with each iteration. We believe this highlights the benefits of our anytime algorithm, giving the user control over the tradeoff between precision and analysis-time.

E. Representative Examples

An inspection of the results indicates that the improvement in precision in SEM-DCIA comes from two fronts. First, it compensates for the price we paid for soundness by considering entire procedures as source of impact. The semantic analysis reduces the impacts for callers and callees transitively. Second, the reduction in impact happens from refactorings that a pure dataflow analysis cannot consider. We show a few interesting patterns that we discovered while applying the tool (for brevity we only describe the change briefly).

Variable Extraction: Figure 5 shows a refactoring to extract a constant to a variable. A non-semantic technique will create

```
-while (*c) {
+for (; *c; c++) {
    ...
    wprintw(window, "%c",
    *c);
    - c++;
}
```

Fig. 6. Change illustrating a loop conversion in mdp commit 00c2ad

Project Name	DCIA			SEM-DCIA ₀				SEM-DCIA ₁				SEM-DCIA _∞			
	min	max	Time	min	max	Red	Time	min	max	Red	Time	min	max	Red	Time
flingfd	64	84	0.94	39	83	20.1%	8.92	14	70	47.3%	9.85	14	70	47.3%	10.44
histo	0	86	2.14	0	75	11.5%	19.43	0	65	28.6%	20.59	0	65	28.6%	24.92
mdp	0	465	28.16	0	330	1.5%	77.71	0	324	3.4%	100.68	0	283	6.5%	173.08
tinyvm	0	344	68.96	0	308	18.5%	158.33	0	298	23.2%	160.35	0	283	43.6%	169.05
theft	184	261	4.48	11	186	61%	38.45	11	185	62%	57.48	11	107	77%	289.75
print_tokens	151	153	2.22	69	137	19.37%	24.02	34	128	28.67%	58.23	34	128	28.67%	102.40
print_tokens2	155	158	1.46	80	129	30.36%	16.08	59	101	44.65%	24.42	55	100	45.66%	97.98
replace	75	195	4.96	74	194	2.08%	35.72	70	194	2.89%	92.72	65	174	9.41%	236.77
schedule	79	115	1.37	7	104	26.35%	13.73	7	87	40.58%	24.15	7	68	70.85%	30.83
space	20	2851	59.45	14	2816	31.87%	798.96	14	2816	36.71%	895.14	n.a.	n.a.	n.a.	timeout
tcas	1	49	0.66	0	49	9.24%	7.94	0	49	9.24%	8.63	0	49	9.24%	9.71
tot_info	103	104	6.39	31	102	18.65%	37.01	24	102	46.26%	74.99	12	77	56.50%	127.61

TABLE III
ANALYSIS RESULTS FOR DIFFERENT LEVELS OF PRECISION. TIME IN SECONDS. (TIMEOUT = 1 HOUR)

```

...
- if (!strend || !strbegin) goto pp_ret;
+ if (!strend || !strbegin) return 0;

    if (!pFile) {
        ...
- goto pp_ret;
+ return 0;
    }
    ...
- pp_ret: return 0;
+ return 0;
}

```

Fig. 7. Change illustrating a goto-elimination refactoring in tinyvm commit 378cc6

impacts in `term_move_to` through the first argument, since it will not be able to find that the value flowing into the first argument is the same in both versions and in all executions. Our SEM-DCIA technique will successfully prove the mutual precondition necessary to show the equality in both versions, and hence cut impacts that would propagate through the first argument.

Loop Refactoring: Figure 6 shows a change from a while loop to a for loop. Input-output equivalence checking would not prevent the impact of the argument `c` to the callee inside the loop (nor would dataflow analysis). Remember that we extract loops as tail recursive procedures.

Control-Flow Equivalence: Figure 7 shows a change to replace a goto with return statements. This is a change in the project tinyvm. The goto statements were all redirecting control-flow to a return statement, so the developer replaced the goto with the target return statement.

VI. RELATED WORK

Our work is closely related to work on change-impact analysis, regression verification, regression test generation, and relational analysis.

Change impact analysis: Change Impact Analysis has been widely explored in static and dynamic program analysis context [10], [32], [34], [45], [48]. Most previous works perform the analysis at a coarse-grain level (classes and types) to retain soundness of analysis [1], [2], [33], [39] which can result in coarse results. JDiff [1] addresses some of the challenges of performing both a diff and computing a mapping between two programs in the context of Java object-oriented programs. Other techniques resort to dynamic information to recover from the overly-conservative dataflow analysis [2], [39]. Our goal is to improve the precision of CIA analysis by making it

change-semantics aware using statically computed equivalence relations without sacrificing soundness.

Regression verification: Regression verification [22], [42] and its implementations [30] aim at proving summary equivalence interprocedurally, but does not help with the CIA directly as shown in § I-A. The work by Bakes et al. [3] improves traditional equivalence checking by pruning paths not impacted by changes. The approach is non-modular (does not summarize callees), bounded (unrolls loops and recursion), and does not seek to improve the underlying change-impact analysis. The technique leverages CIA to avoid performing equivalence checking on non-impacted procedures (computed by standard dataflow analysis). These approaches are useful for equivalence-preserving changes; when the changes are non-equivalent they do not provide meaningful help for reducing code review or testing efforts. Our approach, on the other hand, refines the CIA and can be used in code review and regression testing. Besides, our approach retains modularity and is sound in the presence of loops and recursion. We leverage the product construction in SYMDIFF [31] that has been used for *differential assertion checking* (checking if an assertion fails more often after a change); however this work is limited as it requires the presence of assertions in the program. Our approach can also use other product construction techniques and relational invariant inference techniques as an off-the-shelf solver [7], [8], [11].

Regression testing: Person et al. using change directed symbolic execution to generate regression tests [43]. Our technique can be used to prune the space of statements for which regression tests need to be generated. In addition, there is research on relational verification using a product construction [7]–[9], [40], but most approaches are not automated and do not consider changes across procedure calls.

VII. CONCLUSIONS

In this work, we have formalized and demonstrated how to leverage equivalence relations to improve the precision of change-impact analysis and provide a scalability-precision knob with SEM-DCIA-ANYTIME, which is crucial for applying such analyses to large projects. Our work brings together program verification techniques (namely relational invariant generation) to improve the precision of a core software engineering task, and can go a long way in providing the benefits of deep semantic reasoning to average developers.

REFERENCES

- [1] T. Apiwattanapong, A. Orso, and M. J. Harrold. A differencing algorithm for object-oriented programs. In *Proceedings of the 19th IEEE international conference on Automated software engineering*, pages 2–13. IEEE Computer Society, 2004.
- [2] T. Apiwattanapong, A. Orso, and M. J. Harrold. Efficient and precise dynamic impact analysis using execute-after sequences. In *Proceedings of the 27th international conference on Software engineering*, pages 432–441. ACM, 2005.
- [3] J. Backes, S. Person, N. Rungta, and O. Tkachuk. Regression verification using impact summaries. In *Model Checking Software*, pages 99–116. Springer, 2013.
- [4] M. Barnett, B.-Y. E. Chang, R. DeLine, B. Jacobs, and K. R. M. Leino. Boogie: A modular reusable verifier for object-oriented programs. In *International Symposium on Formal Methods for Components and Objects (FMCO)*, pages 364–387, 2006.
- [5] M. Barnett, K. R. M. Leino, and W. Schulte. The spec# programming system: An overview. In *Construction and analysis of safe, secure, and interoperable smart devices*, pages 49–69. Springer, 2004.
- [6] C. Barrett, A. Stump, and C. Tinelli. The SMT-LIB standard: Version 2.0. In *International Workshop on Satisfiability Modulo Theories (SMT)*, 2010.
- [7] G. Barthe, J. M. Crespo, and C. Kunz. Relational verification using product programs. In *FM 2011: Formal Methods*, pages 200–214. Springer, 2011.
- [8] G. Barthe, J. M. Crespo, and C. Kunz. Beyond 2-safety: Asymmetric product programs for relational program verification. In *Logical Foundations of Computer Science*, pages 29–43. Springer, 2013.
- [9] N. Benton. Simple relational correctness proofs for static analyses and program transformations. In *ACM SIGPLAN Notices*, volume 39, pages 14–25. ACM, 2004.
- [10] H. Cai and R. Santelices. A comprehensive study of the predictive accuracy of dynamic change-impact analysis. *Journal of Systems and Software*, 103:248–265, 2015.
- [11] M. Carbin, D. Kim, S. Misailovic, and M. C. Rinard. Proving acceptability properties of relaxed nondeterministic approximate programs. In *ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, pages 169–180, 2012.
- [12] J. Condit, B. Hackett, S. K. Lahiri, and S. Qadeer. Unifying type checking and property checking for low-level code. In *ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL)*, pages 302–314, 2009.
- [13] Coreutils paste.c commit. <https://github.com/coreutils/coreutils/commit/8297f568690319556cf1424531f71598662b>.
- [14] Coreutils sort.c commit. <https://github.com/coreutils/coreutils/commit/611e7e02bfb889c622d6ed582a9232c746b644>.
- [15] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL)*, pages 238–252, 1977.
- [16] T. L. Dean and M. S. Boddy. An analysis of time-dependent planning. In *AAAI*, volume 88, pages 49–54, 1988.
- [17] J.-R. Falleri, F. Morandat, X. Blanc, M. Martinez, and M. Montperrus. Fine-grained and accurate source code differencing. In *Proceedings of the 29th ACM/IEEE international conference on Automated software engineering*, pages 313–324. ACM, 2014.
- [18] J. Ferrante, K. J. Ottenstein, and J. D. Warren. The program dependence graph and its use in optimization. *ACM Trans. Program. Lang. Syst.*, 9(3):319–349, 1987.
- [19] J. Ferrante, K. J. Ottenstein, and J. D. Warren. The program dependence graph and its use in optimization. *ACM Trans. Program. Lang. Syst.*, 9(3):319–349, 1987.
- [20] C. Flanagan, K. R. M. Leino, M. Lillibridge, G. Nelson, J. B. Saxe, and R. Stata. Extended static checking for java. In *Proceedings of the 2002 ACM SIGPLAN Conference on Programming Language Design and Implementation (PLDI)*, Berlin, Germany, June 17–19, 2002, pages 234–245, 2002.
- [21] M. Gligoric, L. Eloussi, and D. Marinov. Practical regression test selection with dynamic file dependencies. In *Proceedings of the 2015 International Symposium on Software Testing and Analysis*, pages 211–222. ACM, 2015.
- [22] B. Godlin and O. Strichman. Regression verification. In *DAC*, pages 466–471, 2009.
- [23] P. A. Grillet. *Semigroups: an introduction to the structure theory*, volume 193. CRC Press, 1995.
- [24] A. Gyori, S. Lahiri, and N. Partush. Interprocedural semantic change-impact analysis using equivalence relations. Technical report, University of Illinois, August 2016. <http://mir.cs.illinois.edu/gyori/publications/tr-scia.pdf>.
- [25] C. Hawblitzel, M. Kawaguchi, S. K. Lahiri, and H. Rebelo. Towards modularly comparing programs using automated theorem provers. In *International Conference on Automated Deduction (CADE)*, pages 282–299. Springer, 2013.
- [26] M. Hutchins, H. Foster, T. Goradia, and T. Ostrand. Experiments of the effectiveness of dataflow-and controlflow-based test adequacy criteria. In *Proceedings of the 16th international conference on Software engineering*, pages 191–200. IEEE Computer Society Press, 1994.
- [27] J.-M. Jezequel and B. Meyer. Design by contract: The lessons of ariane. *Computer*, 30(1):129–130, 1997.
- [28] M. Kawaguchi, S. Lahiri, and H. Rebelo. Conditional equivalence. Technical report, Microsoft Research, October 2010.
- [29] J. Krinke. Identifying similar code with program dependence graphs. In *Reverse Engineering, 2001. Proceedings. Eighth Working Conference on*, pages 301–309. IEEE, 2001.
- [30] S. K. Lahiri, C. Hawblitzel, M. Kawaguchi, and H. Rebelo. SymDiff: A language-agnostic semantic diff tool for imperative programs. In *International Conference on Computer Aided Verification (CAV)*, pages 712–717, 2012.
- [31] S. K. Lahiri, K. L. McMillan, R. Sharma, and C. Hawblitzel. Differential assertion checking. In *Joint Meeting of the European Software Engineering Conference and ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE)*, pages 345–355, 2013.
- [32] J. Law and G. Rothermel. Whole program path-based dynamic impact analysis. In *Software Engineering, 2003. Proceedings. 25th International Conference on*, pages 308–318. IEEE, 2003.
- [33] W. Le and S. D. Pattison. Patch verification via multiversion interprocedural control flow graphs. In *Proceedings of the 36th International Conference on Software Engineering*, pages 1047–1058. ACM, 2014.
- [34] S. Lehnert. A review of software change impact analysis. *Ilmenau University of Technology, Tech. Rep.*, 2011.
- [35] F. Logozzo, S. Lahiri, M. Fahndrich, and S. Blackshear. Verification modulo versions: Towards usable verification. In *Proceedings of the 35th conference on Programming Languages, Design, and Implementation (PLDI 2014)*. ACM SIGPLAN, June 2014.
- [36] P. D. Marinescu and C. Cadar. make test-zesti: A symbolic execution solution for improving regression testing. In *Proceedings of the 34th International Conference on Software Engineering*, pages 716–726. IEEE Press, 2012.
- [37] S. K. Lahiri, K. L. McMillan, R. Sharma, and A. E. Hassan. The impact of new participation on software quality: A case study of the qt, vtk, and itk projects. In *Proceedings of the 11th Working Conference on Mining Software Repositories*, pages 192–201. ACM, 2014.
- [38] K. L. McMillan. An interpolating theorem prover. In *International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, pages 16–30, 2004.
- [39] A. Orso, T. Apiwattanapong, and M. J. Harrold. Leveraging field data for impact analysis and regression testing. In *ACM SIGSOFT Software Engineering Notes*, volume 28, pages 128–137. ACM, 2003.
- [40] N. Partush and E. Yahav. Abstract semantic differencing for numerical programs. In *Static Analysis - 20th International Symposium, SAS 2013, Seattle, WA, USA, June 20–22, 2013. Proceedings*, pages 238–258, 2013.
- [41] F. Pastore, L. Mariani, A. E. Hyvärinen, G. Fedukovich, N. Sharygina, S. Sehestedt, and A. Muhammad. Verification-aided regression testing. In *Proceedings of the 2014 International Symposium on Software Testing and Analysis*, pages 37–48. ACM, 2014.
- [42] S. Person, M. B. Dwyer, S. Elbaum, and C. S. Păsăreanu. Differential symbolic execution. In *Proceedings of the 16th ACM SIGSOFT International Symposium on Foundations of Software Engineering, SIGSOFT ’08/FSE-16*, pages 226–237, New York, NY, USA, 2008. ACM.
- [43] S. Person, G. Yang, N. Rungta, and S. Khurshid. Directed incremental symbolic execution. In *Proceedings of the 32nd ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’11*, pages 504–515, New York, NY, USA, 2011. ACM.
- [44] Z. Rakamaric and M. Emmi. SMACK: decoupling source language details from verifier implementations. In *Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18–22, 2014. Proceedings*, pages 106–113, 2014.

- [45] X. Ren, F. Shah, F. Tip, B. G. Ryder, and O. Chesley. Chianti: a tool for change impact analysis of java programs. In *ACM Sigplan Notices*, volume 39, pages 432–448. ACM, 2004.
- [46] T. W. Reps, S. Horwitz, and S. Sagiv. Precise interprocedural dataflow analysis via graph reachability. In *Conference Record of POPL’95: 22nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, San Francisco, California, USA, January 23-25, 1995*, pages 49–61, 1995.
- [47] G. Rothermel and M. J. Harrold. A safe, efficient regression test selection technique. *ACM Trans. Softw. Eng. Methodol.*, 6(2):173–210, Apr. 1997.
- [48] B. G. Ryder and F. Tip. Change impact analysis for object-oriented programs. In *Proceedings of the 2001 ACM SIGPLAN-SIGSOFT workshop on Program analysis for software tools and engineering*, pages 46–53. ACM, 2001.
- [49] Wikipedia. Anytime algorithm. https://en.wikipedia.org/wiki/Anytime_algorithm.
- [50] W. Yang. Identifying syntactic differences between two programs. *Software: Practice and Experience*, 21(7):739–755, 1991.
- [51] S. Zilberstein and S. Russell. Approximate reasoning using anytime algorithms. *Kluwer International Series in Engineering and Computer Science*, pages 43–43, 1995.

APPENDIX

Inferring Equalities

In this section, we will describe one particular implementation of PREEQUIV and SUMMARYEQUIV using a product construction in SYMDIFF [31]. We choose the product construction in SYMDIFF because it allows for precise interprocedural differential reasoning over the product of two programs; most other product constructions are limited to single procedures with loops. We start by briefly summarizing the product construction (detailed description is outside the scope of this work) and the inference of equivalences over the product. We ~~next~~ identify restrictions under which these equivalences serve as valid PREEQUIV and SUMMARYEQUIV and augment the product construction to ensure these restrictions.

A. Product Program Construction

Given $Prog^1$ and $Prog^2$, SYMDIFF generates a product program $Prog^{1 \times 2}$ that defines a procedure $f^{1 \times 2}$ for every $f \in Procs$. Loops are already assumed to be extracted as tail-recursive procedures in $Procs$. Figure 8 shows the high-level structure of a product procedure; we will ignore the highlighted statement for now. The input and output formals of $f^{1 \times 2}$ is a concatenation of the inputs and outputs of f^1 and f^2 , suitably renamed. The body of $f^{1 \times 2}$ consists of 3 parts. (1) Line 4 initializes two variables $callHist^i$ that capture the sequence of callsites in a trace through the body of f^i . (2) Lines 5, ..., 10 and Lines 11, ..., 16 contain the bodies of f^1 and f^2 respectively. Any callsite in either f^i (underlined) is instrumented to track (a) the arguments (b) returns and (c) callsite (e.g. line 8). (3) Finally, lines 18, ..., 27 iterate through all pair of callsites in a path in f^1 and f^2 with the same callee (say g), and invokes the product procedure $g^{1 \times 2}$ with corresponding arguments (line 22, and constrains the outputs to be outputs from the respective callsites (line 24). The product procedure $main^{1 \times 2}$ that is the entry procedure for $Prog^{1 \times 2}$ has a precondition to constrain the input arguments $\bar{x}^1$ and $\bar{x}^2$ to start in equal states.

For a procedure f with inputs $\bar{x}$ and outputs $\bar{r}$, a Boolean expression $\phi(\bar{x}) \Rightarrow \psi(\bar{x}, \bar{r})$ is a *summary* for f if for any $(\theta_1, \theta_2) \in \Omega_f$, if θ_1 satisfies ϕ , then θ_2 satisfies ψ . The following theorem relates $f^{1 \times 2}$ with f^1 and f^2 .

Theorem 4 ([31]): If $\phi(\bar{x}^1, \bar{x}^2) \Rightarrow \psi(\bar{x}^1, \bar{x}^2, \bar{r}^1, \bar{r}^2)$ is a summary for a product procedure $f^{1 \times 2}$, then for any $(\theta_1, \theta_2) \in \Omega_f^1$ and $(\theta_3, \theta_4) \in \Omega_f^2$, if $\theta_1 \oplus \theta_3$ satisfies ϕ then $\theta_2 \oplus \theta_4$ satisfies ψ .

B. Equivalence inference

For the product program $Prog^{1 \times 2}$, one can leverage any of the (single program) invariant generation techniques based on abstract-interpretation [15], or interpolants [38] to infer preconditions, postconditions (including two-state postconditions) on $f^{1 \times 2}$. Such invariants are *relational* or *differential* in that they are over the state of two programs $Prog^1$ and $Prog^2$. Equivalences are special classes of such relational invariants and therefore we leverage the invariant generation over $Prog^{1 \times 2}$ to infer PREEQUIV and SUMMARYEQUIV. In

```

1  proc  $f^{1 \times 2}(\bar{x}^1, \bar{x}^2): (\bar{r}^1, \bar{r}^2)$ 
2  {
3      /-- initialize call sequences --
4      var  $callHist^1, callHist^2 := \emptyset, \emptyset;$ 
5      /-- Start of  $f^1$  body
6      ...
7       $n_1: \text{call } \bar{r}^1 := g^1(\bar{e}^1); // \bar{r}^1 \cap RVars(\bar{e}^1) = \emptyset$ 
8           $callHist^1 += (n_1, g, \bar{e}^1, \bar{r}^1); // \text{record call}$ 
9          ...
10     /-- End of  $f^1$  body
11     /-- Start of  $f^2$  body
12     ...
13      $n_2: \text{call } \bar{r}^2 := h^2(\bar{e}^2); // \bar{r}^2 \cap RVars(\bar{e}^2) = \emptyset$ 
14          $callHist^2 += (n_2, g, \bar{e}^2, \bar{r}^2); // \text{record call}$ 
15         ...
16     /-- End of  $f^1$  body
17     /-- constrain calls
18     foreach (var  $g$  in  $Procs$ ) {
19         foreach ( $(n_i, g, \bar{e}^1, \bar{y}^1)$  in  $callHist^1$ ) {
20             foreach ( $(n_j, g, \bar{e}^2, \bar{y}^2)$  in  $callHist^2$ ) {
21                 //replay inputs
22                  $\text{call } \bar{o}^1, \bar{o}^2 := g^{1 \times 2}(\bar{e}^1, \bar{e}^2);$ 
23                 //constrain outputs
24                 assume  $(\bar{o}^1 == \bar{y}^1 \wedge \bar{o}^2 == \bar{y}^2);$ 
25             }
26         }
27     }
28     assert  $callHist^1 ==_{procs} callHist^2;$ 
29 }
30 proc  $main^{1 \times 2}(\bar{x}^1, \bar{x}^2)$ 
31 free requires  $\bar{x}^1 == \bar{x}^2$ 
32 { ... }

```

Fig. 8. Product procedures for inferring equivalences [31].

the next few paragraphs, we identify conditions under which such relational invariants can yield our desired equivalences.

For an output formal $r \in Out_f$, let $SummaryEq(r^1, r^2)$ be the predicate $(\bigwedge_{x \in In_f \text{ and } x \in Deps(r)} x^1 == x^2) \Rightarrow r^1 == r^2$.

Theorem 5: For a product program $Prog^{1 \times 2}$ and a procedure $f^{1 \times 2}$ and output parameter $r \in Out_f$, if $SummaryEq(r^1, r^2)$ is a summary of $f^{1 \times 2}$, then $SUMMARYEQUIV(r, f)$ holds, provided all executions in f^1 and f^2 are terminating.

Theorem 6: For a product program $Prog^{1 \times 2}$ and a procedure $f^{1 \times 2}$ and input parameter $x \in In_f$, if $x^1 == x^2$ is a precondition of $f^{1 \times 2}$, then $PREEQUIV(x, f)$ holds, provided:

- 1) All executions in $Prog^1$ and $Prog^2$ are terminating,
- 2) For any procedure $h \in Procs$, and for any pair of traces $\tau_h^1(\theta)$ and $\tau_h^2(\theta)$ starting in θ , the sequences of procedures calls are identical.

```

void  $main()$ 
{
     $f(1, 2);$ 
}

int  $f(\text{int } x, \text{int } y)$ 
{
    int  $r = x + 2;$ 
     $+ f(r, y);$ 
    return  $r;$ 
}

```

Fig. 9. Example illustrating side conditions for inference.

a) *Example:* Figure 9 illustrates the need for the side conditions when working with the product program. The change in f adds an infinite loop in $Prog^2$ by calling itself. Here $Deps(r) = \{x\}$ for f in both programs. Since f does not terminate, neither does $f^{1 \times 2}$. Hence, $(x^1 == x^1) \Rightarrow r^1 == r^2$ is

a vacuous summary of $f^{1 \times 2}$, although $SUMMARYEQUIV(y, f)$ does not hold. On the other hand, since f does not invoke itself in $Prog^1$, no calls to $f^{1 \times 2}$ are made from within $f^{1 \times 2}$. This means that $x^1 == x^2$ is a precondition for $f^{1 \times 2}$, although $PREEQUIV(x, f)$ does not hold.

Termination is a desirable property for a program and can be verified separately (e.g. establishing ranking functions for loops and recursive procedures. In fact, we can relax the need for every execution to terminate by only requiring *equitermination* (an input θ leads to a terminating execution in $Prog^1$ if and only if $Prog^2$ terminates on the same input) [25]; this is an area of future work. For this work, we (unsoundly) assume that the changes do not introduce or fix non-termination.

Satisfying the second side condition for Theorem 6, unfortunately, cannot be reduced to a standard program verification task. We informally describe a modification to the product construction (Figure 8) that provides a sufficient check for ensuring that call sequences from a common input are identical. We reuse the history variables $callHist^i$ that capture the call sequences along any path in either procedures (lines 8 and 14). We add an assertion (highlighted in line 28) that checks if call sequence through f^1 is identical to f^2 (assuming equitermination); the equality $==_{procs}$ compares the two sequences after projecting on to the second element (procedure name) of each tuple. If *every* product procedure satisfies this assertion, starting from only the assumption that arguments to *main* are identical, then one can ensure the second side condition for Theorem 6. Our implementation currently does not handle this modification.

APPENDIX

Anytime Algorithm

It is worth highlighting that small bounds for k may not always suffice to achieve precise results, as illustrated by the example in Figure 10. For this example $Procs^\Delta = \{main\}$; The set $Procs'$ grows as follows across iterations: $\{main\}$, $\{f_{2n}, main, f_1\}$, $\{f_{2n-1}, f_{2n}, main, f_1, f_2\}$, and so on. It is not hard to see that until $Procs' = Procs$, $InferEquivs$ will not be able to infer that the $PREEQUIV(x, f_1)$ holds.

```

void  $main(\text{int } x)$ 
{
    -  $f_1(x);$ 
    +  $f_1(x+0);$ 
}
void  $f_1(\text{int } x)$ 
{
     $f_2(x+1);$ 
}
...

void  $f_{2n-1}(\text{int } x)$ 
{
     $f_{2n}(x+2n-1);$ 
}
void  $f_{2n}(\text{int } x)$ 
{
    if ( $x > 0$ )
         $f_1(x-1);$ 
}

```

Fig. 10. Pathological example that needs large k