A New Foundation For Control-Dependence and Slicing for Modern Program Structures*

Venkatesh Ranganath[1], Torben Amtoft[1], Anindya Banerjee[1], Matthew B. Dwyer[2], and John Hatcliff[1]

[1] Department of Computing and Information Sciences, Kansas State University **
[2] Department of Computer Science and Engineering, University of Nebraska, Lincoln ***

Abstract. The notion of control dependence underlies many program analysis and transformation techniques used in numerous applications. Despite wide application, existing definitions and approaches to calculating control dependence are difficult to apply seamlessly to modern program structures. Such programs structures make substantial use of exception processing and increasingly support reactive systems designed to run indefinitely.
This paper revisits foundational issues surrounding control dependence and develops definitions and algorithms for computing control dependence that can be directly applied to modern program structures. In the context of slicing reactive systems, the paper proposes a notion of slicing correctness based on weak bisimulation and proves that the definition of control dependence generates slices that conform to this notion of correctness. Finally, a variety of properties show that the new definitions conservatively extend classic definitions. These new definitions and algorithms for control dependence form the basis of a publicly available program slicer that has been implemented for full Java.

1 Introduction

The notion of control-dependence underlies many program analysis and transformation techniques used in numerous applications including program slicing applied for program understanding [1], debugging [2], and optimizations, partial evaluation [3], compiler optimizations [4] such as global scheduling, loop fusion, code motion etc. Intuitively, a program statement n_1 is control-dependent on a statement n_2, if n_2 (typically, a conditional statement) controls whether or not n_1 will be executed or bypassed during an execution of the program.

While existing definitions and approaches to calculating control dependence and slicing are widely applied and have been used in the current form for well over 20 years, there are several aspects of these definitions that prevent them from being applied smoothly to modern program structures which rely significantly on exception processing and increasingly support reactive systems which are designed to run indefinitely.

(I.) Classic definitions of control dependence are stated in terms of program control-flow graphs (CFGs) in which the CFG has a unique end node – they do not apply directly to program CFGs with (a) multiple end nodes or with (b) no end node. Restriction (a) means that existing definitions cannot be applied directly to programs/methods with multiple exit points – a restriction that would be violated by any method that raises exceptions or includes multiple returns. Restriction (b) means that existing definitions

* This work was supported in part by the U.S. Army Research Office (DAAD190110564), by DARPA/IXO's PCES program (AFRL Contract F33615-00-C-3044), by NSF (CCR-0306607) by Lockheed Martin, and by Intel Corporation.
** Manhattan KS, 66506, USA. {rvprasad,tamtoft,ab,hatcliff}@cis.ksu.edu
*** Lincoln NE 68588-0115, USA. {dwyer}@cse.unl.edu

cannot be applied directly to reactive programs or system models with control loops that are designed to run indefinitely.

Restriction (a) is usually addressed by performing a pre-processing step that transforms a CFG with multiple end nodes into a CFG with a single end node by adding a new designated end node to the CFG and inserting arcs from all original exit states to the new end node [5, 1] Restriction (b) can also be addressed in a similar fashion by, e.g., selecting a single node within the CFG to represent the end node. This case is more problematic than the pre-processing for Restriction (a) because the criteria for selecting end nodes that lead to the desired control dependence relation between program nodes is often unclear. This is particularly true in threads such as event-handlers which have no explicit shut-down methods, but are "shut down" by killing the thread (thus, there is nothing in the thread's control flow to indicate an exit point).

(II.) Existing definitions of slicing correctness either apply to programs with terminating execution traces, or they often fail to state whether or not the slicing transformation preserves the termination behavior of the program being sliced. Thus these definitions cannot be applied to reactive programs that are designed to execute indefinitely. Such programs are used in numerous modern applications such as event-processing modules in GUI systems, web services, distributed real time systems with autonomous components, e.g. data sensors, etc.

Despite the difficulties, it appears that researchers and practitioners do continue to apply slicing transformations to programs that fail to satisfy the restrictions above. However, in reality the pre-processing transformations related to issue **(I)** introduce extra overhead into the entire transformation pipeline, clutter up program transformation and visualization facilities, necessitate the use/maintenance of mappings from the transformed CFGs back to the original CFGs, and introduce extraneous structure with ad-hoc justifications that all down-stream tools/transformations must interpret and build on in a consistent manner. Moreover, regarding issue **(II)** it will be infeasible to continue to ignore issues of termination as slicing is increasingly applied in high-assurance applications such as reducing models for verification [6] and for reasoning about security issues where it is crucial that liveness/non-termination properties be preserved.

Working on a larger project on slicing concurrent Java programs, we have found it necessary to revisit basic issues surrounding control dependence and have sought to develop definitions that can be directly applied to modern program structures such as those found in reactive systems. In this paper, we propose and justify the usefulness and correctness of simple definitions of control dependence that overcome the problematic aspects of the classic definitions described above. The specific contributions of this paper are as follows.

- We propose new definitions of control dependence that are simple to state and easy to calculate and that work directly on control-flow graphs that may have no end nodes or non-unique end nodes, thus avoiding troublesome pre-processing CFG transformations (Section 4).
- We prove that these definitions applied to reducible CFGs yield slices that are correct according to generalized notions of slicing correctness based on a form of weak-bisimulation that is appropriate for programs with infinite execution traces (Section 5.1).
- We clarify the relationship between our new definitions and classic definitions by showing that our new definitions represent a form of "conservative extension" of classic definitions: when our new definitions are applied to CFGs that conform to the restriction of a single end node, our definitions correspond to classic definitions – they do not introduce any additional dependences nor do they omit any dependences. (Section 5.1).

– We discuss the intuitions behind algorithms for computing control dependence (according to the new definitions) to justify that control dependence is computable in polynomial time (Section 6).

Expanded discussions, definitions and full proofs appear in the companion technical report [7] which can be found on the project web site [8].

The proposed notions of control dependence described in this paper have been implemented in Indus-Kaveri [8] – our publicly available open-source Eclipse-based Java slicer that works on full Java 1.4 and has been applied to code bases of up to 10,000 lines of Java application code (< 80K bytecodes) excluding library code. Besides its application as a stand-alone program visualization, debugging, and code transformation tool, our slicer is being used in the next generation of our Bandera tool set for model-checking concurrent Java systems.

2 Basic Definitions

2.1 Control Flow Graphs

When dealing with foundational issues of control dependence, researchers often cast their work in terms of a simple imperative language phrased in terms of control flow graphs. We follow that practice here and base our presentation on a definition of control-flow graph adapted from Ball and Horwitz [9].

Definition 1 (Control Flow Graphs).
A control-flow graph $G = (N, E, n_0)$ is a labeled directed graph in which

- *N is a set of nodes that represent commands in program,*
- *the set of N is partitioned into two subsets N^S, N^P, where N^S are* statement nodes *with each $n_s \in N^S$ having at most one successor, where N^P are* predicate nodes *with each $n_p \in N^P$ having two successors, and $N^E \subseteq N^S$ contains all nodes of N^S that have no successors, i.e., N^E contains all end nodes of G,*
- *E is a set of labeled edges that represent the control flow between graph nodes where each $n_p \in N^P$ has two outgoing edges labeled* T *and* F *respectively, and each $n_s \in (N^S - N^E)$ has an outgoing edge labeled* A *(representing* A*lways taken),*
- *the start node n_0 has no incoming edges and all nodes in N are reachable from n_0.*

We will display the labels on CFG edges only when necessary for the current exposition.

As stated earlier, existing presentations of slicing require that each CFG G satisfies the *unique end node property*: there is exactly one element in $N^E = \{n_e\}$ and n_e is reachable from all other nodes of G. The definition above *does not* require this property of CFGs, but we will sometimes consider CFGs with the unique end node property in our comparisons to previous work.

To relate a CFG with the program that it represents, we use the function *code* to map a CFG node n to the code for the program statement that corresponds to that node. Specifically, for $n_s \in N^S$, $code(n_s)$ yields the code for an assignment statement, and for $n_p \in N^P$, $code(n_p)$ the code for the test of a conditional statement (the labels on the edges for n_p allow one to determine the nodes for the true and false branches of the conditional). The function *def* maps each node to the set of variables defined (*i.e.*, assigned to) at that node (always a singleton or empty set), and *ref* maps each node to the set of variables referenced at that node.

A CFG *path* π from n_i to n_k is a sequence of nodes $n_i, n_{i+1}, \ldots, n_k$ such for every consecutive pair of nodes (n_j, n_{j+1}) in the path there is an edge from n_j to n_{j+1}. A path between nodes n_i and n_k can also be denoted as $[n_i..n_k]$. When the meaning is clear from the context, we will use π to denote the set of nodes contained in π and we write $n \in \pi$

when n occurs in the sequence π. Path π is *non-trivial* if it contains at least two nodes. A path is *maximal* if it is infinite or if it terminates in an end node.

The following definitions describe relationships between graph nodes and the distinguished start and end nodes [10]. Node n *dominates* node m in G (written $dom(n, m)$) if every path from the start node s to m passes through n (note that this makes the dominates relation reflexive). Node n *post-dominates* node m in G (written $post\text{-}dom(n, m)$) if every path from node m to the end node e passes through n. Node n *strictly post-dominates* node m in G if $post\text{-}dom(n, m)$ and $n \neq m$. Node n is the *immediate post-dominator* of node m if $n \neq m$ and n is the first post-dominator on every path from m to the end node e. Node n *strongly post-dominates* node m in G if n post-dominates m and there is an integer $k \geq 1$ such that every path from node m of length $\geq k$ passes through n [1]. The difference between strong post-domination and the simple definition of post-domination above is that even though node n occurs on every path from m to e (and thus n post-dominates m), it may be the case that there is a loop in the CFG between m and n that admits an infinite path beginning at m that never encounters n. Strong post-domination rules out the possibility of such loops between m and n – thus, it is sensitive to the possibility of non-termination along paths from m to n. Note that domination relations are well-defined but post-domination relationships are not well-defined for graphs that do not have the unique end node property.

A CFG G is *reducible* if E can be partitioned into disjoint sets E_f (the *forward* edge set) and E_b (the *back* edge set) such that (N, E_f) forms a DAG in which each node can be reached from the entry node n_0 and for all edges $e \in E_b$, the target of e dominates the source of e. All "well-structured" programs give rise to reducible control-flow graphs, including Java programs. Our definitions and most of our correctness results apply to irreducible CFGs as well, but our bi-simulation-based correctness of slicing result only holds for reducible graphs since bi-simulation requires ordering properties that can only be guaranteed on reducible graphs.

2.2 Program Execution

The execution semantics of program CFGs is phrased in terms of transitions on program states (n, σ) where n is a CFG node and σ is a store mapping the corresponding program's variables to values. A series of transitions gives an *execution trace* through p's statement-level control flow graph. It is important to note that when execution is in state (n_i, σ_i), the code at node n_i has not yet been executed. Intuitively, the code at n_i is executed on the transition from (n_i, σ_i) to successor state (n_{i+1}, σ_{i+1}). Execution begins at the state node n_0, and the execution of each node possibly updates the store and transfers control to an appropriate successor node. Execution of a node $n_e \in N^E$ produces a final state (halt, σ) where the control point is indicated by a special label halt – this indicates a normal termination of program execution. The presentation of slicing in the next section involves arbitrary finite and infinite non-empty sequences of states written $\Pi = s_1, s_2, \ldots$. For a set of variables V, we write $\sigma_1 =_V \sigma_2$ when for all $x \in V$, $\sigma_1(x) = \sigma_2(x)$.

2.3 Notions of Dependence and Slicing

A *program slice* consists of the parts of a program p that (potentially) affect the variable values that are referenced at some program points of interest [11]. Traditionally, the program "points of interest" are called the *slicing criterion*. A slicing criterion C for a program p is a non-empty set of nodes $\{n_1, \ldots, n_k\}$ where each n_i is a node in p's CFG.

The definitions below recall the two basic notions of dependence that appear in slicing of sequential programs: *data dependence* and *control dependence* [11].

Data dependence captures the notion that a variable reference is dependent upon any variable definition that "reaches" the reference.

Definition 2 (data dependence). *Node n is* data-dependent *on m (written $m \stackrel{dd}{\rightarrow} n$ – the arrow pointing in the direction of data flow) if there is a variable v such that*

1. *there exists a non-trivial path π in p's CFG from m to n such that for every node $m' \in \pi - \{m, n\}$, $v \notin def(m')$, and*
2. *$v \in def(m) \cap ref(n)$.*

Control dependence information identifies the conditionals that may affect execution of a node in the slice. Intuitively, node n is control-dependent on a predicate node m if m directly determines whether n is executed or "bypassed".

Definition 3 (control dependence). *Node n is* control-dependent *on m in program p (written $m \stackrel{cd}{\rightarrow} n$) if*

1. *there exists a non-trivial path π from m to n in p's CFG such that every node $m' \in \pi - \{m, n\}$ is post-dominated by n, and*
2. *m is not strictly post-dominated by n.*

For a node n to be control-dependent on predicate m, there must be two paths that connect m with the unique end node e such that one contains n and the other does not. There are several slightly different notions of control-dependence appearing in the literature, and we will consider several of these variants and relations between them in the rest of the paper. At present, we simply note that the above definition is standard and widely used (e.g., see [10]).

We write $m \stackrel{d}{\rightarrow} n$ when either $m \stackrel{dd}{\rightarrow} n$ or $m \stackrel{cd}{\rightarrow} n$. Constructing a program slice proceeds by finding the set of CFG nodes S_C (called the *slice set*) from which the nodes in C are reachable via $\stackrel{d}{\rightarrow}$.

Definition 4 (slice set). *Let C be a slicing criterion for program p. Then the slice set S_C of p with respect to C is defined as follows:*

$$S_C \;=\; \{m \;\mid\; \exists n \,.\, n \in C \;\; and \;\; m \stackrel{d}{\rightarrow}^{*} n\}.$$

The notion of slicing described above is referred to as "backward static slicing" because the algorithm starts at the criterion nodes and looks backward through the program's control-flow graph to find other program statements that influence the execution at the criterion nodes. In this paper we consider only backward slices, but our definitions of control dependence can also be applied we computing forward slices.

In many cases in the slicing literature, the desired correspondence between the source program and the slice is not formalized because the emphasis is often on applications rather than foundations, and this also leads to subtle differences between presentations. When a notion of "correct slice" is given, it is often stated using the notion of *projection* [12]. Informally, given an arbitrary trace Π of p and an analogous trace Π_s of p_s, p_s is a correct slice of p if projecting out the nodes in criterion C (and the variables referenced at those nodes) for both Π and Π_s yields identical state sequences. We will consider slicing correctness requirements in greater detail in Section 5.1.

3 Assessment of Existing Definitions

3.1 Variations in Existing Control Dependence Definitions

Although the definition of control dependence that we stated in Section 2 is widely used, there are a number of (sometimes subtle) variations appearing in the literature. One

dimension of variation is whether the particular definition captures only *direct* control dependence or also admits *indirect* control dependences. For example, using the definition of control dependence in Definition 3, for Fig. 1 (a), we can conclude that $a \stackrel{cd}{\rightarrow} f$ and $f \stackrel{cd}{\rightarrow} g$ however $a \stackrel{cd}{\rightarrow} g$ does not hold because g does not post-dominate f. The fact that a and g are indirectly related (a does play a role in determining if g is executed or bypassed) is not captured in the definition of control dependence itself but in the transitive closure used in the slice set construction (Definition 4). However, some definitions of control dependence [1] incorporate this notion of transitivity directly into the definition itself as we will illustrate later.

Another dimension of variation is whether the particular definition is sensitive to non-termination or not. Consider Fig. 1 (a) where node c represents a post-test that controls a loop – which may be infinite (one cannot tell by simply looking at the CFG). According to Definition 3, $a \stackrel{cd}{\rightarrow} d$ but $c \stackrel{cd}{\rightarrow} d$ does not hold (because d post-dominates c) even though c may determine whether d executes or never gets to execute due to an infinite loop that postpones d forever. Thus, Definition 3 is *non-termination insensitive*.

We now further illustrate these dimensions by recalling definitions of strong and weak control dependence given by Podgurski and Clarke [1] and used in a number of works including the study of control dependence by Bilardi and Pingali [13].

Definition 5 (Podgurski-Clarke Control Dependence).

- *n_2 is strongly control dependent on n_1 ($n_1 \stackrel{PC-scd}{\rightarrow} n_2$) if there is a path from n_1 to n_2 that does not contain the immediate post dominator of n_1.*
- *n_2 is weakly control dependent on n_1 ($n_1 \stackrel{PC-wcd}{\rightarrow} n_2$) if n_2 strongly post dominates n_1', a successor of n_1, but does not strongly post dominate n_1'', another successor of n_1.*

The notion of strong control dependence above roughly corresponds to Definition 3, but it captures indirect control dependence whereas Definition 3 captures only direct control dependence. For example, in Fig. 1, in contrast to Definition 3 we have $a \stackrel{PC-scd}{\rightarrow} g$ because there is a path afg which does not contain the immediate post-dominator of a. However, one can show that when used in the context of Definition 4 (which computes the transitive closure of dependences), the two definitions give rise to the same slices.

The notion of weak control dependence above subsumes the notion of strong control dependence ($n_1 \stackrel{PC-scd}{\rightarrow} n_2$ implies $n_1 \stackrel{PC-wcd}{\rightarrow^*} n_2$) and it captures weaker dependences between nodes induced by non-termination, that is, it is non-termination sensitive. Note that for Fig. 1 (a), $c \stackrel{PC-wcd}{\rightarrow} d$ because d does not strongly post-dominate b: the presence of the loop controlled by c guarantees that there does not exist a k such that every path from node b of length $\geq k$ passes through d.

In assessing the above variants of control dependence in the context of program slicing, it is important to note that slicing based on Definition 3 or the strong control dependence above can transform a non-terminating program into a terminating one (i.e., non-termination is not preserved in the slice). In Fig. 1 (a), assume that the loop controlled by c is an infinite loop. Using the slice criterion $C = \{d\}$ would include a but not b and c (we assume no data dependence between d and b or c) if the slicing is based on strong control dependence. Thus, in the sliced program, one would be able to observe an execution of d, but such an observation is not possible in the original program because execution diverges before d is reached. In contrast, the difference between direct and indirect statements of control dependence seem to largely technical stylistic decision in how the definitions are stated.

Very few works consider the non-termination sensitive notion of weak control dependence above. We conjecture that there are at least two reasons for this. First, although it bears the qualifier "weak", weak control dependence is actually a stronger relation

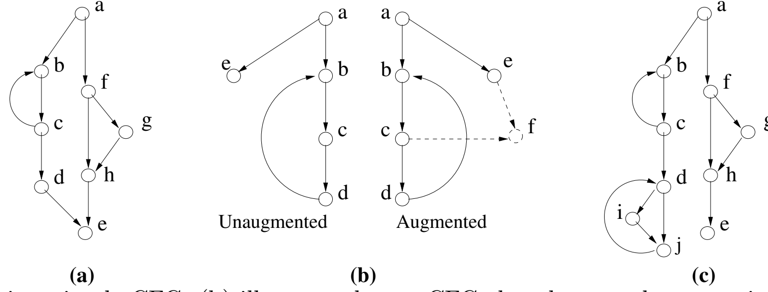

Fig. 1. (a) is a simple CFG. (b) illustrates how a CFG that does not have a unique exit node reachable from all nodes can be augmented to have unique exit node reachable from all nodes. (c) is a CFG with multiple control sinks of different sorts.

(relating more nodes) and will thus include more nodes in the slice. Second, many applications of slicing focus on debugging and program visualization and understanding, and in these applications having slices that preserve non-termination is less important than having smaller slices. However, slicing is increasingly used in security applications and as a model-reduction technique for software model checking. In these applications, it is quite important to consider variants of control dependence that preserve non-termination properties since failure to do so could allow inferences to be made that compromise security policies, for instance invalidate checks of liveness properties [6].

3.2 Unique End node restriction on CFG

All definitions of control dependences that we are aware of require that CFGs satisfy the unique end node requirement – but many software systems fail to satisfy this property. Existing works simply require that CFGs have this property, or they suggest that CFGs can be augmented to achieve this property, e.g., using the following steps: (1) insert a new node e into the CFG, (2) add an edge from each exit node (other than e) to e, (3) pick an arbitrary node n in each non-terminating loop and add an edge from n to e. In our experience, such augmentations complicate the system being analyzed in several ways. If the augmentation is non-destructive, a new CFG is generated which costs time and memory. If the augmentation is destructive, this may clash with the requirements of other clients of the CFG, thus necessitating the reversal of the augmentation before subsequent analyses can proceed. In addition, having multiple end nodes (e.g., an exceptional exit and a regular return) flow into a single new end node causes semantically different information to flow together.

Many systems have threads where the main control loop has no exit – the loop is "exited" by simply killing the thread. For example, in Xt library, most applications create widgets, register callbacks, and call `XtAppMainLoop()` to enter an infinite loop that manages the dispatching of events to the widgets in the application. In PalmOS, applications are designed such that they start upon receiving a start code, execute a loop, and terminate upon receiving a stop code. However, the application may choose to ignore the stop code once it starts, and hence, not terminate except when it is explicitly killed. In such cases, a node in the loop must be picked as the loop exit node for the purpose of augmenting the CFG. However, this can disrupt the control dependence calculations. In Fig. 1 (b), we would intuitively expect e,b,c, and d to be control dependent on a in the unaugmented CFG. However, $a \overset{PC-wcd}{\rightarrow} \{e, b, c\}$ and $c \overset{PC-wcd}{\rightarrow} \{b, c, d, f\}$ in the augmented CFG. It is trivial to prune dependences involving f. However, there are new dependences $c \overset{PC-wcd}{\rightarrow} \{b, c, d\}$ which did not exist in the unaugmented CFG. Although a suggestion to delete any dependence on c may work for the given CFG, it fails if there exists a node g that is a successor of c and a predecessor of d. Also, $a \overset{PC-wcd}{\rightarrow} d$ exists in the

unaugmented CFG but not in the augmented CFG, and it is not obvious how to recover this information.

We address these issues head-on by considering alternate definitions of control-dependence that do not impose the unique end-node description.

4 New Dependence Definitions

In previous definitions, a control dependence relationship where n_j is dependent on n_i is specified by considering paths from n_i and n_j to a unique CFG end node – essentially n_i and the end node delimit the path segments that are considered. Since we aim for definitions that apply when CFGs do not have an end node or have more than one end node, we aim to instead specify that n_j is control dependent on n_i by focusing on paths between n_i and n_j. Specifically, we focus on path segments that are delimited by n_i *at both ends* – intuitively corresponding to the situation in a reactive program where instead of reaching an end node, a program's behavior begins to repeat itself by returning again to n_i. At a high level, the intuition remains the same as in, e.g., Definition 3 – executing one branch of n_i always leads to n_j, whereas executing another branch of n_i can cause n_j to be bypassed. The additional constraints that are added (e.g., n_j always occurs before any occurrence of n_i) limits the region in which n_j is seen or bypassed to segments leading up to the next occurrence of n_i – ensuring that n_i is indeed *controlling* n_j. The definition below considers maximal paths (which includes infinite paths) and thus is sensitive to non-termination.

Definition 6 ($n_i \stackrel{ntscd}{\rightarrow} n_j$). *In a CFG, n_j is* ***(directly) non-termination sensitive control dependent*** *on node n_i if n_i has at least two successors, n_k and n_l,*

(1) for all maximal paths from n_k, n_j always occurs and it occurs before any occurrence of n_i. occurrence of n_i
(2) there exists a maximal path from n_l on which either n_j does not occur, or n_j is strictly preceded by n_i.

We supplement a traditional presentation of dependence definitions with definitions given as formulae in computation tree logic (CTL) [14]. CTL is a logic for describing the structure of sets of paths in a graph, making it a natural language for expressing control dependences. Informally, CTL includes two path quantifiers, E and A, which define that a path from a given node with a given structure exists or that all paths from that node have the given structure. The structure of a path is defined using one of five modal operators (we refer to a node satisfying ϕ as a ϕ-node): $\mathsf{X}\phi$ states that the successor node is a ϕ-node, $\mathsf{F}\phi$ states the existence of a ϕ-node, $\mathsf{G}\phi$ states that a path consists entirely of ϕ-nodes, $\phi\mathsf{U}\psi$ states the existence of a ψ-node and that the path leading up to that node consists of ϕ-nodes, finally, the $\phi\mathsf{W}\psi$ operator is a variation on U that relaxes the requirement that a ψ-node exist. In a CTL formula path quantifiers and modal operators occur in pairs, e.g., $\mathsf{AF}\phi$ says on all paths from a node a ϕ node occurs. A formal definition of CTL can be found in [14].

The following CTL formula captures the definition of control dependence above.

$$n_i \stackrel{ntscd}{\rightarrow} n_j = (G, n_i) \models \mathsf{EX}(\mathsf{A}[\neg n_i \mathsf{U} n_j]) \wedge \mathsf{EX}(\mathsf{E}[\neg n_j \mathsf{W}(\neg n_j \wedge n_i)]).$$

Here, $(G, n_i) \models$ expresses the fact that the CTL formula is checked against the graph G at node n_i. The two conjuncts are essentially a direct transliteration of the natural language above.

We have formulated the definition above to apply to *execution traces* instead of CFG paths. In this setting one needs to bound relevant segments by n_i as discussed above. However, when working on CFG paths, the definition conditions can actually be

simplified to read as follows: (1) *for all maximal paths from n_k, n_j always occurs*, and (2) *there exists a maximal path from n_l on which n_j does not occur*. The corresponding CTL formula is

$$n_i \overset{ntscd}{\rightarrow} n_j = (G, n_i) \models \mathsf{EX}(\mathsf{AF}(n_j) \wedge \mathsf{EX}(\mathsf{EG}(\neg n_j)).$$

See [7] for the proof that these two definitions are equivalent on CFGs.

To see that this definition is non-termination sensitive, note that $c \overset{ntscd}{\rightarrow} d$ in Fig. 1 (a) since there exists a maximal path (an infinite loop between b and c) where d never occurs. Moreover, the definition corresponds to our intuition in Section 3.2 in that, in Fig. 1 (b unaugmented) $a \overset{ntscd}{\rightarrow} e$ because there is an infinite loop through b, c, d and $a \overset{ntscd}{\rightarrow} \{b, c, d\}$ because there is maximal path ending in e that does not contain b, c, or d. In Fig. 1 (c), note that $d \overset{ntscd}{\rightarrow} i$ because there is an infinite path from j (cycle on j,d) on which i does not occur.

We now turn to constructing a non-termination insensitive version of control dependence. The definition above considered all paths leading out of a conditional. Now, we need to limit the reasoning to finite paths that reach a terminal region of the graph. To handle this in the context of CFGs that do not have the unique end-node property, we generalize the concept of *end node* to *control sink* – a set of nodes such that each node in the set is reachable from every other node in the set and there is no path leading out of the set. More precisely, a *control sink* κ is a set of CFG nodes that form a strongly connected component such that for each $n \in \kappa$ each successor of n is also in κ. It is trivial to see that each end node forms a control sink and each loop without any exit edges in the graph forms a control sink. For example, $\{e\}$ and $\{b, c, d\}$ are control sinks in Fig. 1 (b unaugmented), and $\{e\}$ and $\{d, i, j\}$ are control sinks in Fig. 1 (c). *c-sink* denotes a set-valued function on nodes such that $c\text{-}sink(n) = S$ where if n belongs to a control sink then S is set of nodes representing that sink, otherwise $S = \emptyset$.

Existing definitions of non-termination insensitive control dependence rely on reasoning about paths from the conditional to the end node. We generalize this to reason about paths from a conditional to control sinks. The set of *sink-bounded paths from n_k* (denoted $SinkPaths(n_k)$) contains all π such that π is a path from n_k to a node n_s such that n_s belongs to a control sink.

Definition 7 ($n_i \overset{nticd}{\rightarrow} n_j$). *In a CFG, n_j is **(directly) non-termination insensitively control dependent** on n_i if n_i has at least 2 successors, n_k and n_l,*

(1) for all paths $\pi \in SinkPaths(n_k)$, $n_j \in \pi$.
(2) there exists a path $\pi \in SinkPaths(n_l)$ such that $n_j \notin \pi$ and if π leads to a control sink κ, $n_j \notin \kappa$.

This definition is expressed in CTL as

$$n_i \overset{nticd}{\rightarrow} n_j \;=\; (G, n_i) \models \mathsf{EX}(\hat{\mathsf{A}}\mathsf{F}(n_j)) \wedge \mathsf{EX}(\hat{\mathsf{E}}[\neg n_j \mathsf{U}(c\text{-}sink? \wedge n_j \notin c\text{-}sink)])$$

where $\hat{\mathsf{A}}$ and $\hat{\mathsf{E}}$ represent quantification over sink-bounded paths only. *c-sink?* evaluates to *true* only if the current node belongs to a control sink and *c-sink* returns the sink set associated with the current node.

To see that this definition is non-termination insensitive, note that $c \overset{nticd}{\not\rightarrow} d$ in Fig. 1 (a) since there does exists path from b to a control sink ($\{e\}$ is the only control sink) that does not contain d. Again, in Fig. 1 (b unaugmented) $a \overset{nticd}{\rightarrow} e$ because there path from b to the control sink $\{b, c, d\}$ and neither the path nor the sink contain e, and $a \overset{nticd}{\rightarrow} \{b, c, d\}$ because there is path ending in control sink $\{e\}$ that does not contain b, c, or d. It is interesting to note that in Fig. 1 (c), our definition concludes that $d \overset{nticd}{\not\rightarrow} i$

because although there is a trivial path from d to the control sink $\{d, i, j\}$, i belongs to that control sink. This is because our definition inherently captures a form of fairness – since the backedge from j guarantees that d will be executed an infinite number of times, the only way to avoid executing i would be to branch to d on every cycle. The consequence of this property is that even though there may be control structures inside of a control sink, they will not give rise to any control dependences. In applications where one desires to detect such dependences, one would apply the definition to control sinks in isolation with back edges removed.

4.1 Properties of the Dependence Relations

We begin by showing that the new definitions of control dependence conservatively extend classic definitions: when we consider our definitions in the original setting with CFGs with unique end nodes, the definitions coincide with the classic definitions. In addition, direct non-termination insensitive control dependence (Definition 7) implies the *transitive closure* of direct non-termination sensitive control dependence.

Theorem 1 (Coincidence Properties). *For all CFGs with with the unique end node property, and for all nodes $n_i, n_j \in N$,*

(1) $n_i \neq n_j$ and $n_i \stackrel{cd}{\rightarrow} n_j$ implies $n_i \stackrel{nticd}{\rightarrow} n_j$
(2) $n_i \stackrel{nticd}{\rightarrow} n_j$ implies $n_i \stackrel{cd}{\rightarrow} n_j$
(3) $n_i \stackrel{PC-wcd}{\rightarrow} n_j$ iff $n_i \stackrel{ntscd}{\rightarrow} n_j$
(4) For all CFGs, for all nodes $n_i, n_j \in N : n_i \stackrel{nticd}{\rightarrow} n_j$ implies $n_i \stackrel{ntscd^*}{\rightarrow} n_j$

For the correctness (bisimulation-based) proof in Section 5.1, we shall need a few results about slice sets (members of which are "observable"). A crucial property is that the first observable on any path will be encountered sooner or later on all other paths:

Lemma 1. *Assume the node set Ξ is closed under termination sensitive control dependency, and that $n_0 \notin \Xi$. Assume that there is a path π from n_0 to n_1, with $n_1 \in \Xi$ but for all $n \in \pi$ with $n \neq n_1$, $n \notin \Xi$. Then all maximal paths from n_0 will contain n_1.*

Proof. Assume, in order to arrive at a contradiction, that there exists a maximal path from n_0 that does not contain n_1. We define a predicate Q, such that $Q(n)$ holds iff there exists a maximal path from n that does not contain n_1. By our assumption, $Q(n_0)$ holds; clearly, $Q(n_1)$ does not hold. Therefore, π can be written as $[n_0..n_2 n_3..n_1]$ where $Q(n_2)$ holds but $Q(n_3)$ does not hold (that is, there is an edge from n_2 to n_3; note that n_2 may equal n_0 and that n_3 may equal n_1 but we know that $n_1 \neq n_2$).

We shall show that $n_2 \stackrel{ntscd}{\rightarrow} n_1$; then from $n_1 \in \Xi$ we from Ξ being closed under $\stackrel{ntscd}{\rightarrow}$ get $n_2 \in \Xi$ which contradicts n_1 being the only node in π which is also in Ξ.

Note that since $Q(n_2)$ holds, there exists a maximal path starting at n_2 not containing n_1; that path has to have at least two elements (since n_2 has an outgoing edge) and the second element cannot be n_3 (as $Q(n_3)$ does not hold). Therefore, the second element is some node n_4 with $n_3 \neq n_4$, and there exists a maximal path from n_4 which does not contain n_1. Our final obligation is to prove that all maximal paths from n_3 contain n_1, which follows since $Q(n_3)$ does not hold.

In a similar way we can show:

Lemma 2. *Assume Ξ is closed under $\stackrel{nticd}{\rightarrow}$, and that $n_0 \notin \Xi$. Assume that there is a path π from n_0 to n_1, with $n_1 \in \Xi$ but for all $n \in \pi$ with $n \neq n_1$, $n \notin \Xi$. Then all sink-bounded paths from n_0 will contain n_1.*

As a consequence we have the following result, giving conditions to preclude the existence of infinite un-observable paths:

Lemma 3. *Assume that $n_0 \notin \Xi$, but that there is a path π starting at n_0 which contains a node in Ξ.*

- *If Ξ is closed under termination insensitive control dependency, then all sink bounded paths starting at n_0 will reach Ξ.*
- *If Ξ is also closed under termination sensitive control dependency, then all maximal paths starting at n_0 will reach Ξ.*

We are now ready for the main result, stating that from a given node there is a unique first observable (for this, we need the CFG to be reducible, as can be seen by the counterexample where from n_0 there are edges to n_1 and n_2 between which there is a cycle).

Theorem 2. *Assume that $n_0 \notin \Xi$, that $n_1, n_2 \in \Xi$, and that there are paths $\pi_1 = [n_0..n_1]$ and $\pi_2 = [n_0..n_2]$ such that on both paths, all nodes except the last do not belong to Ξ.*

If Ξ is closed under termination insensitive control dependency (a weaker requirement than being closed under termination sensitive control dependency), and if the CFG is reducible, then $n_1 = n_2$.

Proof. Clearly, we can extend π_1 and π_2 into sink-bounded paths π'_1 and π'_2. By Lemma 2, we infer that π'_2 contains n_1, and that π'_1 contains n_2. If $n_1 \neq n_2$, this implies that n_1 is reachable from n_2, and vice versa, while both being reachable from n_0, something which cannot happen in a reducible graph.

5 Slicing

We now describe how to slice a (reducible) CFG G wrt. a slice set S_C, the smallest set containing C which is closed under data dependence $\stackrel{dd}{\rightarrow}$ and also under some kind of control dependence: at least we must require it is closed under $\stackrel{nticd}{\rightarrow}$, but a stronger correctness property (Sect. 5.1) holds if it is also closed under $\stackrel{ntscd}{\rightarrow}$.

The result of slicing is a program with the same CFG as the original one, but with the code map $code_1$ replaced by $code_2$. Here $code_2(n) = code_1(n)$ for $n \in S_C$; for $n \notin S_C$ then

- if n is a statement node then $code_2(n)$ is the statement `skip`;
- if n is a predicate node then $code_2(n)$ is `cskip`, the semantics of which is that it non-deterministically chooses one of its successors.

The above definition is conceptually simple, so as to facilitate the correctness proofs. Of course, one would want to do some post-processing, like eliminating `skip` commands and eliminating `cskip` commands where the two successor nodes are equal; we shall not address this issue further but remark that most such transformations are trivially meaning preserving.

5.1 Correctness Properties

The main intuition behind our notion of slicing correctness is that the nodes in a slicing criteria C represent "observations" that one is making about a CFG G under consideration. Specifically, for a $n \in C$, one can observe that n has been executed and also observe the values of any variables referenced at n. Execution of nodes not in C correspond to *silent moves* or non-observable actions. The slicing transformation should preserve the behavior of the program with respect to C-observations, but parts of the program that

are irrelevant with respect to computing C observations can be "sliced away". The slice set S_C built according to Definition 4 represents the nodes that are relevant for maintaining the observations C. Thus, to prove the correctness of slicing we will establish the stronger result that G will have the same S_C observations wrt. the original code map $code_1$ as wrt. the sliced code map $code_2$, and this will imply that they have the same C observations.

The discussion above suggests that appropriate notions of correctness for slicing reactive programs can be derived from the notion of weak bisimulation found in concurrency theory, where a transition may include a number of τ-moves [15]. In our setting, we shall consider transitions that do one or more steps before arriving at a node in the slice set.

Definition 8. *For $i = 1, 2$ we write $s \stackrel{i}{\longmapsto} s'$ to denote that wrt. code map $code_i$, the program state s rewrites in one step to s'.*

For $i = 1, 2$ we write $s_0 \stackrel{i}{\Longrightarrow} s$ if there exists $s_1 \ldots s_k$ ($k \geq 1$) with $s_k = s$ such that (with each $s_j = (n_j, \sigma_j)$)

- *for all $j \in \{1 \ldots k\}$ we have $s_{j-1} \stackrel{i}{\longmapsto} s_j$;*
- *$n_k \in S_C$ but for all $j \in \{1 \ldots k-1\}$; $n_j \notin S_C$.*

Definition 9. *A binary relation $\mathcal{S}$ on program states is a bisimulation if whenever $(s_1, s_2) \in \mathcal{S}$ then*

(a) if $s_1 \stackrel{1}{\Longrightarrow} s_1'$ then there exists a s_2' such that $s_2 \stackrel{2}{\Longrightarrow} s_2'$ and $(s_1', s_2') \in \mathcal{S}$, and
(b) if $s_2 \stackrel{2}{\Longrightarrow} s_2'$ then there exists a s_1' such that $s_1 \stackrel{1}{\Longrightarrow} s_1'$ and $(s_1', s_2') \in \mathcal{S}$.

If instead of (b) we only have (c) below, we say that $\mathcal{S}$ is a quasi-bisimulation.

(c) if $s_2 \stackrel{2}{\Longrightarrow} s_2'$ then either $s_1 \stackrel{1}{\not\Longrightarrow}$ or there exists a s_1' such that $s_1 \stackrel{1}{\Longrightarrow} s_1'$ and $(s_1', s_2') \in \mathcal{S}$.

For each node n in G, we define $relv(n)$, the set of relevant variables at n, by stipulating that $x \in relv(n)$ if there exists a node $n_k \in S_C$ and a path π from n to n_k such that $x \in refs(n_k)$, but $x \notin defs(n_j)$ for all nodes n_j occurring before n_k in π.

The above is well-defined in that it does not matter whether we use $code_1$ or $code_2$, as it is easy to see that the value of $relv(n)$ is not influenced by the content of nodes not in S_C, since that set is closed under $\stackrel{dd}{\rightarrow}$. (Also, the closedness properties of S_C are not affected by using $code_2$ rather than $code_1$.)

We are now ready to state the correctness theorem:

Theorem 3. *Let the relation $\mathcal{S}_0$ be given by $(n_1, \sigma_1)\, \mathcal{S}_0\, (n_2, \sigma_2)$ iff $n_1 = n_2$ and $\sigma_1 =_{relv(n_1)} \sigma_2$. Then (if G is reducible)*

- *$\mathcal{S}_0$ is a quasi-bisimulation;*
- *$\mathcal{S}_0$ is even a bisimulation if S_C is closed under $\stackrel{ntscd}{\rightarrow}$.*

Proof. (Sketch.) We must consider transitions of the form $(n, \sigma_i) \stackrel{i}{\Longrightarrow} (n', \sigma_i')$; that is we have $(n, \sigma_i) \stackrel{i}{\longmapsto} (n'', \sigma_i'')$ and either $n'' = n'$ or $(n'', \sigma_i'') \stackrel{i}{\Longrightarrow} (n', \sigma_i')$.

With $j = 3 - i$, our general goal is to simulate the above transition wrt. $code_j$. For three cases, listed below, we find σ_j'' such that $(n, \sigma_j) \stackrel{j}{\longmapsto} (n'', \sigma_j'')$ with $\sigma_i'' =_{relv(n'')} \sigma_j''$: then we are done if $n'' = n$; otherwise we apply inductive reasoning.

$n \in S_C$ Here $\sigma_i =_{ref(n)} \sigma_j$. Therefore, if n is a predicate, the same branch will be taken; if n is a statement, the stores will be updated with the same value.

$n \notin S_C$ is a statement Here $code_2(n) = \texttt{skip}$, and the claim follows since the value stored by $code_1(n)$ will not belong to $relv(n'')$ (as S_C is closed under $\stackrel{dd}{\rightarrow}$).

$n \notin S_C$ is a predicate, $i = 1$ Then $code_2(n) = \texttt{cskip}$, and the claim is trivial.

We are left with the interesting case where $n \notin S_C$ is a predicate, $i = 2$. Two subcases:

- S_C is closed under $\stackrel{ntscd}{\rightarrow}$; we must show (b) of Definition 9. But Lemma 3 tells us that there exists n_1, σ_1' such that $(n, \sigma_1) \stackrel{1}{\Longrightarrow} (n_1, \sigma_1')$, where $n_1 = n'$ by Theorem 2. For $x \in relv(n')$, we have to show that $\sigma_1'(x) = \sigma_2'(x)$, which follows since such variables cannot be modified along the way (again since S_C is closed under $\stackrel{dd}{\rightarrow}$).
- otherwise, we only have to show (c) of Definition 9, so assume that there exists n_1, σ_1' such that $(n, \sigma_1) \stackrel{1}{\Longrightarrow} (n_1, \sigma_1')$. By Theorem 2 we infer that $n_1 = n'$, and we proceed as in the previous case.

6 Non-Termination Sensitive Control Dependence ($\stackrel{ntscd}{\longrightarrow}$) Algorithm

Control dependences are calculated using a symbolic data-flow analysis. Fundamentally, control dependences are determined by reasoning about properties of sets of CFG paths; those sets are represented symbolically in our algorithm. Specifically, for each node n with more than one successor in G, the set of paths starting at n that begin with $n \rightarrow m$ is represented by t_{nm}. The algorithm propogates these symbolic values to collect the effects of particular control flow choices at program points in the CFG. For each node p in the CFG a set of symbolic values, S_{pn}, is stored for each node $n \neq p$ in the CFG that has more than one successor; these sets record the set of paths that originate from n. The algorithm preserves the invariant that if t_{nm} is in S_{pn} then all paths from n starting with $n \rightarrow m$ contain node p. A complete description of the algorithm, its correctness and its complexity are given in [7]; in the rest of this section we provide an overview of its main processing steps and its adaptation to computing other forms of dependence.

Let T_n denote the outdegree of n and $condNodes(G)$ denote the set of nodes with outdegree greater than one.

The algorithm is initialized such that, for each node $n \in condNodes(G)$, t_{nm} is inserted into S_{mn} for each successor m of n and m is marked for processing. The algorithm then proceeds by executing the following three steps for each marked node m; it terminates when there are no longer any marked nodes.

1. For each node $n \in condNodes(G) \backslash m$, if $|S_{mn}| = T_n$ then, for each node $p \in condNodes(G) \backslash m$, all symbols from S_{nn_4} are inserted into S_{mp}. This captures the property that if all non-terminal paths or terminal paths that end in exit nodes from every successor of n contains m, then these paths will also contain p.
2. Depending on the number of successors of m, one of the following actions is performed if any S_{nq} was changed.

   $|succ(m)| = 1$ Let p be the successor of m. For each node n such that $S_{pn} \backslash S_{mn} \neq \emptyset$, insert S_{mn} into S_{pn} and add p ito the worklist. This captures the property that all non-terminal paths or terminal paths that end in exit nodes that contain m will also contain n.

   $|succ(m)| > 1$ For each node n, if $|S_{nm}| = T_m$ then n is marked for processing. This captures the requirement that any path information change at m needs to be considered at each node n that will occur on all non-terminal paths or terminal paths that end in exit nodes starting from m.
3. Unmark m.

When there are no more marked nodes, all-path reachability information for every pair of nodes, n and m (with outdegree greater than one), in the graph is available in S_{nm}. The presence of a token t_{mp} in S_{nm} indicates that all non-terminal paths or terminal paths that end in exit nodes starting with the edge $m \rightarrow p$ contain n. So, if $|S_{nm}| > 0 \wedge |S_{nm}| \neq T_m$ then, by Definition 6, it can be inferred that n *is directly control dependent on* m. On the other hand, if $|S_{nm}| > 0$ and $|S_{nm}| = T_m$ then, by Definition 6, it can be inferred that n *is not directly control dependent on* m.

The proposed algorithms have a worst-case asymptotic complexity of $O(|N|^3 \times K)$ where K is the sum of the outdegree of all nodes with more than one successor in the CFG. Linear time algorithms to calculate control dependence have been proposed in the literature [1]. These algorithms, however, rely on augmentation of the CFG. The practical cost of this augmentation varies with the specific algorithm and control dependence being calculated. Our experience with an implementation of our general algorithms in a program slicer for full Java suggests that, despite its complexity bound, it can be scaled to programs with tens-of-thousands of lines of code and still return results in a matter of seconds. We suspect that this is due in part to the elimination of the need for augmenting CFGs in our approach.

7 Related Work

Fifteen years ago, control dependence was rigorously explored by Podgurski and Clarke in [1]. Since, then there has been a variety of work related to calculation and application of control dependence in the setting of CFGs that satisfy the unique end node property.

In the realm of calculating control dependence, Bilardi et.al [13] proposed new concepts related to control dependence along with algorithms based on these concepts to efficiently calculate weak control dependence. In [16], Johnson proposed an algorithm that could be used to calculate control dependence in time linear in the number of edges. In comparison, in this paper we sketch a feasible algorithm in a more general setting.

In the context of slicing, Horwitz, Reps, and Binkley [17] presented what has now become the standard approach to inter-procedural slicing via dependence graphs. Recently, Allen and Horwitz [18] extended previous work on slicing to handle exception-based inter-procedural control flow. In this work, they handle CFG's with two end nodes (one for normal return and one for exceptional return) but it is unclear how this affects the control dependence captured by dependence graph. In comparison, we have shown program slicing is feasible with unaugmented CFGs, and the extended version of the paper describes in greater detail how our definitions interact with interprocedural slicing.

For relevant work on slicing correctness, [19], Horwitz et.al. use a semantics based multi-layered approach to reason about the correctness of slicing in the realm of data dependence. In [9], Ball et.al used program point specific history based approach to prove the correctness of slicing for arbitrary control flow. We build off of that work to consider arbitrary control flow with out the unique end-node restriction. Their correctness property is a weaker property than bi-simulation – it does not require ordering to be maintained between observable nodes if there is no dependence between these nodes – and it holds for irreducible CFGs. Even though our definitions apply to irreducible graphs, we need to extra structure of reducible graphs to achieve the stronger correctness property. We are currently investigating if we can establish their correctness property using our control dependence definitions on irreducible graphs.

In [5], Hatcliff et.al. presented notions of dependence for concurrent CFGs, and proposed a notion of bi-simulation as the correctness property. Millett and Teitelbaum [20] study static slicing of Promela (the model description language for the model-checker SPIN) and its application to model checking, simulation, and protocol understanding, but they do not formalize a notion of correct slice nor do they discuss issues related to

preserving non-termination and liveness properties. Krinke [21] considers static slicing of multi-threaded programs with shared variables, and focuses on issues associated with inter-thread data dependence but does not consider non-termination sensitive forms of control dependence.

8 Conclusion

The notion of control dependence is used in myriad of applications, and researchers and tool builders increasing seek to apply it to modern software systems and high-assurance applications – even though the control flow structure and semantic behavior of these systems does not mesh well with the requirements of existing control dependence dependences. In this paper, we have proposed conceptually simple definitions of control dependence that (a) can be applied directly to the structure of modern software thus avoiding unsystematic preprocessing transformations that introduce overhead, conceptual complexity, and sometimes dubious semantic interpretations, and (b) provide a solid semantic foundation for applying control dependence to reactive systems where program executions may be non-terminating.

We have rigorously justified these definitions by detailed proofs, by expressing them in temporal logic which provides an unambiguous definition and allows them to be mechanically checked/debugged against examples using automated verification tools, by showing their relationship to existing definitions, and by implementing and experimenting with them in a publicly available slicer for full Java. In addition, we have provided algorithms for computing these new control dependence relations, and argued that any additional cost in computing these relations is negligible when one considers the cost and ill-effects of preprocessing steps required for previous definitions. Thus, we believe that there are many benefits for widely applying these definitions in static analysis tools.

In ongoing work, we continue to explore the foundations of static and dynamically calculating dependences for concurrent Java programs for slicing, program verification, and security applications. In particular, we are exploring the relationship between dependences extracted from execution traces and dependences extracted from control-flow graphs in an effort to systematically a justify a comprehensive set of dependence notions for the rich features found in concurrent Java programs.

References

1. Podgurski, A., Clarke, L.: A formal model of program dependences and its implications for software testing, debugging, and maintenance. IEEE Transactions on Software Engineering **16** (1990) 965–979
2. Francel, M.A., Rugaber, S.: The relationship of slicing and debugging to program understanding. In: Proceedings of the 7th IEEE International Workshop on Program Comprehension (IWPC'99. (1999) 106–113
3. Anderson, L.O.: Program Analysis and Specialization for the C Programming Languages. PhD thesis, DIKU, University of Copenhagen, DIKU, University of Copenhagen, Universit et sparken 1, DK-2100, Copenhagen ∅, Denmark. (1994)
4. Ferrante, J., Ottenstein, K.J., Warren, J.O.: The program dependence graph and it's use in optimization. ACM Transaction of Programming Languages and Systems **9** (1987) 319–349
5. Hatcliff, J., Corbett, J.C., Dwyer, M.B., Sokolowski, S., Zheng, H.: A formal study of slicing for multi-threaded programs with JVM concurrency primitives. In: Proceedings on the 1999 International Symposium on Static Analysis (SAS'99). Lecture Notes in Computer Science (1999)
6. Hatcliff, J., Dwyer, M.B., Zheng, H.: Slicing software for model construction. Journal of Higher-order and Symbolic Computation **13** (2000) 315–353 A special issue containing selected papers from the 1999 ACM SIGPLAN Workshop on Partial Evaluation and Program Manipulation.

7. Ranganath, V.P., Amtoft, T., Banerjee, A., Dwyer, M.B., Hatcliff, J.: A new foundation for control-dependence and slicing for modern program structures. Technical Report 8, Kansas State University (2004) This is available at http://projects.cis.ksu.edu/docman/admin/index.php?editdoc=1&docid=95&group_id=12.
8. SAnToS Laboratory: Indus, a toolkit to customize and adapt Java programs. (This software is available at http://indus.projects.cis.ksu.edu)
9. Ball, T., Horwitz, S.: Slicing programs with arbitrary control-flow. In: Proceedings of the First International Workshop on Automated and Algorithmic Debugging. Volume 749 of Lecture Notes in Computer Science., Springer-Verlag (1993) 206–222
10. Muchnick, S.S.: Advanced Compiler Design & Implementation. Morgan Kaufmann Publishers. Inc., San Francisco, California, USA (1997)
11. Tip, F.: A survey of program slicing techniques. Journal of programming languages **3** (1995) 121–189
12. Weiser, M.: Program slicing. IEEE Transactions on Software Engineering **10** (1984) 352–357
13. Bilardi, G., Pingali, K.: A framework for generalized control dependences. In: Proceedings of the ACM SIGPLAN 1996 conference on Programming language design and implementation, Philadelphia, Pennsylvania, United States, ACM, ACM Press New York, NY, USA (1996) 291–300
14. Edmund M. Clarke, J., Grumberg, O., Peled, D.A.: Model Checking. The MIT Press (1999)
15. Milner, R.: Communication and Concurrency. Prentice Hall, Prentice Hall Europe, Campus 400, Marylands Avenue, Hemel Hempstead, Hertfordshire, HP2 7EZ (1989) ISBN: 0-13-115007-3.
16. Johnson, R., Pingali, K.: Dependence-based program analysis. In: SIGPLAN Conference on Programming Language Design and Implementation. (1993) 78–89
17. Horwitz, S., Reps, T., Binkley, D.: Interprocedural slicing using dependence graphs. ACM Transactions on Programming Language and Systems (1990)
18. Allen, M., Horwitz, S.: Slicing java programs that throw and catch exceptions. In: Procedings of the 2003 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-Based Program Manipulation (PEPM '03), ACM (2003) 44–54
19. Horwitz, S., Pfeiffer, P., Reps, T.W.: Dependence analysis for pointer variables. In: Proceedings of the ACM SIGPLAN '89 Conference on Programming Language Design and Implementation (PLDI'89), ACM (1989) 28–40
20. I.Millett, L., Teitelbaum, T.: Slicing Promela and its applications to model checking, simulation, and protocol understanding. In: Proceedings of the 4th International SPIN Workshop. (1998)
21. Krinke, J.: Static slicing of threaded programs. In: Proc. ACM SIGPLAN/SIGFSOFT Workshop on Program Analysis for Software Tools and Engineering (PASTE'98). (1998) 35–42

A Details

A.1 Algorithm to calculate Non-Termination Sensitive Control Dependence

Proof of correctness We show that phase (1) and (2) of the algorithm are correct by proving that the following loop invariant holds for the outer loop in each phase.

> $t_{n_1 n_2} \in S_{n_3 n_1}$ *implies each non-trivial loop-free segment* $\pi(n_2, n_1?)$ *of each maximal path* $\pi = [n_2..n_1?]$ *contains* n_3.

At the beginning of phase (1), each token set $T_{n_3 n_1}$ is empty. Hence, the invariant is trivially established. In the loops at line 9 and 10, for each immediate successor node n_2 of each conditional node n_1, $t_{n_1 n_2}$ is injected into $T_{n_2 n_1}$. This trivially preserves the invariant at the end of the loop as $n_3 = n_2$ occurs on all segments starting n_2. The loop will terminate as the number of nodes in the graph is finite.

Now the reasoning about phase (2).

```
Non-Termination-Sensitive-Control-Dependence(G)
 1  G(N, E, n0, N^E) :  a control flow graph.
 2  S[|N|, |N|] :  a matrix of sets where S[n1, n2] represents S_{n1 n2}.
 3  T[|N|] :  a sequence of integers where T[n1] denotes T_{n1}.
 4  CD[|N|] :  a sequence of sets.
 5  workbag :  a set of nodes.
 6
 7   # (1) Initialize
 8  workbag ← ∅
 9  for  each n1 in condNodes(G)
10  do succs = succs(n1, G)
11      for  each n2 in succs
12      do workbag ← workbag ∪ {n2}
13          S[n2, n1] ← {t_{n1 n2}}
14
15   # (2) Calculate all-path reachability
16  while workbag ≠ ∅
17  do flag ← false
18      n3 ← remove(workbag)
19      for  each n1 in condNodes(G)\n3
20      do if |S[n3, n1]| = T[n1]
21              then for  each n4 in condNodes(G)\n3
22                      do if S[n1, n4]\S[n3, n4] ≠ ∅
23                              then S[n3, n4] ← S[n3, n4] ∪ S[n1, n4]
24                                      flag = true
25
26      if flag and |succs(n3, G)| = 1
27        then n5 ← select(succs(n3, G))
28                for n4 in condNodes(G)
29                do if S[n5, n4]\S[n3, n4] ≠ ∅
30                        then S[n5, n4] ← S[n5, n4] ∪ S[n3, n4]
31                                workbag ← workbag ∪ {n5}
32        else  if flag and |succs(n3, G)| > 1
33                  then for  each n4 in N
34                          do if |S[n4, n3]| = T[n3]
35                                  then workbag ← workbag ∪ {n4}
36
37   # (3) Calculate non-termination sensitive control dependence
38  for  each n3 in N
39  do for  each n1 in condNodes(G)
40      do if |S[n4, n3]| > 0 and |S[n3, n1]| ≠ T[n1]
41              then CD[n3] ← CD[n3] ∪ {n1}
42
43  return CD
```

Fig. 2. The algorithm to calculate non-termination sensitive control dependence.

Initialization At the beginning of phase (2), the invariant holds as it held at the termination of phase (1).

Maintenance The loops at line 19-23 and at line 28-30 ensure that if all non-trivial loop-free $(n_1, n_1?)$ segments from n_1 contains n_3 then all non-trivial loop-free $(n_4, n_4?)$ segments that contain n_1 will contain n_3.[3] Hence, the loop invariant is maintained.

Termination In each iteration, either the size of a token set increases at least by one or remains the same. Hence, eventually the size of the token sets will stabilize (not increase) preventing additions of elements to the workbag at lines 31 and 35 (by not setting *flag* to *true* in the conditional at 22). Hence, the loop at line 16 will terminate. Upon termination, as the size of token sets remain unchanged, it should be the case that each set reached it's maximal size.

Hence, after phase (1) and (2), for each successor n_2 of $n_1 \in condNodes(G)$, $t_{n_1 n_2} \in S_{n_3 n_1}$ only if each non-trivial loop-free segment $\pi(n_2, n_1?)$ of each maximal path $\pi = [n_2..n_1?]$ contains n_3.

In phase (3), direct control dependence is calculated based on the available reachability information. The termination of this phase is obvious by the finiteness of the nodes and edges of the graph.

Complexity analysis In phase (1), for every node with multiple successors in the CFG, each of its successors is processed. Hence, it leads to a worst-case asymptotic complexity of $O(|E|)$ for phase (1). In phase (3), for each node, every node in the CFG is processed leading to a worst-case asymptotic complexity of $O(|N|^2)$ for this phase.

In phase (2), the loop at line 16 iterates till the size of the token sets represented by S stabilizes. The maximum size of a token set $S[n_1, n_2]$ is given by $T[n_2]$ which is equal to the outdegree of n_2. In each iteration, either the size of a token set increases at least by one or remains the same. In the former case, it contributes an iteration. As the size of the token sets $S[n_1, n_2]$ is bound, all token sets of $S[n_1]$ will stabilize in a total of $\sum T[i]$ or less iterations. The loops in line 19 and 21 contribute $O(|condNodes(G)|^2) \simeq O(|N|^2)$ to each such iteration. Hence, the worst-case complexity of phase (2) will be $O(|N|^3 \times \sum T[i] \times \lg(|N|))$ by factoring in the complexity $O(\lg |N|)$ of set operations.

By combining the above information, the worst-case complexity to calculate indirect control dependence via phase 1, 2, and 3 will be $O(|E| + |N|^3 \times \sum T[i] \times \lg |N| + |N|^2)$. However, as $O(|N|^3 \times \sum T[i] \times \lg |N|)$ dominates $O(|N|^2)$ and $O(|E|)$, the complexity will be $O(|N|^3 \times \sum T[i] \times \lg |N|$ when $\sum T[i] \times \lg |N| > 1$. It will be $O(|N|^2 + |E|))$ when $\sum T[i] = 0$.

As in practice $|condNodes(G)|^2 \approx |N|$, the complexity in the case where $\sum T[i] \times \lg |N| > 1$ will reduce to $O(|N|^2 \times \sum T[i] \times \lg |N|)$.

[3] Suppose there is a non-trivial loop-free $(n_4, n_4?)$ segment that contains n_1 but not n_3. This implies there is a non-trivial loop-free segment from n_1 to n_4 to n_1 that does not contain n_3, hence, leading to a contradiction.