

Precise Interprocedural Dependence Analysis of Parallel Programs [★]

Markus Müller-Olm

Universität Dortmund, FB Informatik, LS 5, 44221 Dortmund, Germany ¹

Abstract

It is known that interprocedural detection of copy constants and elimination of faint code in parallel programs are undecidable problems, if base statements are assumed to execute atomically. We show that these problems become decidable, if this assumption is abandoned. So, the (unrealistic) idealization from program verification “atomic execution of base statements” introduced in order to simplify matters, actually increases the difficulty of these problems from the point of view of program analysis: amazingly these problems become more tractable if we adopt a less idealized, more realistic model of execution.

We introduce an effective abstract domain of *antichains of dependence traces* that allows us to perform a precise interprocedural dependence analysis in (non-atomically executing) parallel programs. The main idea is to trace sequences of dependences exhibited successively by program executions. We define operations on antichains of dependence traces and show that they precisely abstract the corresponding operations on sets of non-atomic program executions. Using these operations, we can analyze dependences by means of an abstract interpretation of constraint systems that characterize sets of program executions of interest. The result of the dependence analysis can in turn be used to detect all copy constants and to eliminate faint code.

While the run-time of the algorithms is exponential in the number of program variables, it is *polynomial in the program size*. Hence, they are polynomial-time algorithms if the number of program variables is bounded. In order to justify their overall exponential run-time, we show that both detection of copy constants and elimination of faint code are intractable (NP-hard) even when the atomic execution idealization is abandoned. This holds already for parallel programs without loops or procedures.

Key words: program analysis, concurrency, dependence, atomicity assumption

Contents

1	Introduction	3
2	Parallel Flow Graphs	6
2.1	Parallel Flow Graphs	6
2.2	Operational Semantics	7
2.3	Atomic Runs	9
2.4	The Run Sets of Ultimate Interest	10
2.5	The Constraint Systems	11
2.6	Discussion	22
3	Non-Atomic Execution	23
3.1	Modeling Non-Atomic Execution by Virtual Variables	25
3.2	A Motivating Example	27
3.3	The Domain of Non-Atomic Run Sets	28
3.4	Discussion	31
4	Dependence Traces	31
4.1	Transparency and Dependences	32
4.2	Dependence Traces	32
4.3	Implication Order	35
4.4	Subsumption Order	36
4.5	A Lattice of Antichains	37
4.6	Short Dependence Traces	40
4.7	The Abstract Domain	43

* The research reported here was partially supported by the RTD project IST-1999-20527 “DAEDALUS” of the European FP5 programme.

Email address: mmo@ls5.cs.uni-dortmund.de (Markus Müller-Olm).

URL: <http://ls5-www.cs.uni-dortmund.de/~mmo> (Markus Müller-Olm).

¹ Current affiliation: FernUniversität in Hagen, FB Informatik, LG PI 5, Universitätsstrasse 1, 58097 Hagen, Germany.

4.8	Pre-Operator	46
4.9	Post-Operator	48
4.10	Sequential Composition	49
4.11	Interleaving	51
4.12	Base Edges	62
4.13	Run-Time	62
4.14	Discussion	64
5	Detecting Copy Constants and Eliminating Faint Code	64
5.1	Copy Constant Detection	65
5.2	Faint Code Elimination	67
5.3	Run-Time	70
6	Intractability	71
6.1	The SAT-Reduction	72
7	Conclusion	74
	Acknowledgements	75
	References	76

1 Introduction

Automatic analysis of parallel programs is known as a notoriously hard problem. A well-known obstacle is the so-called *state-explosion problem*: the number of (control) states of a parallel program grows exponentially with the number of parallel components. Therefore, most practical flow analysis algorithms of concurrent programs conservatively approximate the effects arising from interference of different threads in order to achieve efficiency. An excellent survey on practical research towards analysis of concurrent programs with many references is provided by Rinard [1]. In contrast to this research, we are interested in analyses of parallel programs that are *exact (or precise)* except of the common abstraction of guarded branching to non-deterministic branching that is well-known from analysis of sequential programs.

Surprisingly, certain basic but important dataflow analysis problems can be solved precisely and efficiently for programs with a fork/join kind of parallelism. Corresponding results have been achieved either by generalizing the fixpoint computation techniques common in classic dataflow analysis of sequential programs [2–4] or by automata-theoretic techniques [5,6]. The most far-reaching result is due to Seidl and Steffen [4] who show that all so-called *gen/kill problems* can be solved interprocedurally in fork/join parallel programs efficiently and precisely. This comprises the important class of *bit-vector analyses*, e.g., live/dead-variables analysis, available-expressions analysis, and reaching-definitions analysis [7].

In view of these results it is interesting to ask whether there are other dataflow problems that can precisely be solved for parallel programs. Natural candidates are problems related to transitive variable dependences, like detection of copy constants [8], elimination of faint code [9], and program slicing [10,11]. For sequential languages these problems give rise to simple distributive dataflow frameworks and may be seen as representatives of the next level of difficulty beyond gen/kill problems.

In [12,13] we show that all the problems mentioned in the previous paragraph are undecidable in parallel programs with procedures (parallel interprocedural analysis). Moreover, we show that these problems are PSPACE-complete in case that there are no procedure calls (parallel intraprocedural analysis), and still (co-)NP-complete if also loops are abandoned (parallel acyclic analysis). Unlike previous undecidability results for parallel languages obtained by Bouajjani and Habermehl [14] and Ramalingam [15] these results are independent of explicit synchronization mechanisms.² At first glance this seems to imply that precise program analysis of parallel languages beyond gen/kill problems is hopeless.

There is, however, an assumption in this work that is not that innocent as it may seem: the assumption that base statements of the parallel programs (e.g. assignment statements) execute as atomic steps. While this idealized assumption is not uncommon in the literature, it is hardly realistic in multi-processor environments where a number of concurrently executing processors access a shared memory, because assignments are broken into smaller instructions prior to execution.

Surprisingly, the reductions of [12] break down when the atomic execution assumption for assignment statements is abandoned. Without assuming atomic execution of assignments the subtle game of re-initialization of variables that is crucial for putting the reductions to work can no longer be played.

² Bouajjani and Habermehl who show undecidability of LTL model-checking for parallel languages without synchronization primitives use the LTL formula to synchronize the runs of two parallel threads that simulate a two-counter machine.

In this paper we show that interprocedural detection of copy constants and interprocedural faint code elimination become indeed decidable (in exponential time) if the atomic execution assumption is abandoned. More generally, we show how to do *precise interprocedural analysis of variable dependences in parallel programs*; here precise means precise with respect to non-atomic program executions. So, the (unrealistic) idealization from program verification “atomic execution of assignment statements” intended to simplify matters actually increases the difficulty of these problems from the program analysis point of view: amazingly these problems become more tractable if we adopt a less idealized, more realistic view of execution. This opens up new potential for analysis of parallel programs.

The paper is organized as follows. In Section 2 we define parallel flow graphs as our model of concurrent programs. We furnish them with an operational semantics and define constraint systems that characterize various sets of program executions of interest. For the moment, we still assume atomic execution of base statements.

In Section 3 we explain, why atomic execution of base statements is an unrealistic assumption in a multi-processor environment. In order to capture the semantics of non-atomic execution we re-interpret the operations and constants used in the constraint systems of Section 2. The idea is to break base statements into atomic actions of smaller granularity and to use an interleaving semantics on these atomic actions. The solution of the constraint systems with respect to this new interpretation is taken as the semantic reference point for analysis of parallel programs when the atomic execution assumption is abandoned.

In Section 4 we introduce a domain of *antichains of dependence traces*. We define operations on this domain and show that these operations precisely abstract the corresponding operations on sets of non-atomic program executions. Thus, we can perform precise interprocedural analysis of variable dependences by solving the constraint systems developed in Section 3 over the dependence traces domain. This information can in turn be used to detect copy constants and eliminate faint code. Corresponding algorithms are developed in Section 5. While the run-time of these algorithms is exponential in the number of program variables, it is *polynomial in the program size*. Hence, they are polynomial-time algorithms if the number of program variables is bounded. In order to justify their overall exponential run-time, we show in Section 6 that both detection of copy constants and elimination of faint code are intractable (NP-hard) even when the atomic execution idealization is abandoned. This holds already for parallel programs without loops or procedures.

Throughout this paper we assume that the reader is familiar with the basic techniques and results from the theory of computational complexity [16,17],

program analysis [18–20,7], and abstract interpretation [21,22].

2 Parallel Flow Graphs

In this section, we introduce a flow graph model for parallel programs (cf. [4,2,23]). Edges in the flow graph are annotated with a base statement, a call of a single procedure, or a parallel call of two procedures. As base statements we allow assignment statements and the do-nothing statement **skip**. We assume that branching is non-deterministic, a common abstraction in flow analysis.

2.1 Parallel Flow Graphs

Let X be a finite set of (global) *program variables* and \mathbf{Expr} a set of expressions (or terms) over X . The precise nature of expressions is immaterial for the moment; we only need that each variable $x \in X$ is also an expression: $X \subseteq \mathbf{Expr}$, and that we can determine for an expression $t \in \mathbf{Expr}$ the set of variables occurring in t , $\mathbf{var}(t) \subseteq X$. Let $\mathbf{Stmt} := \{x := t \mid x \in X, t \in \mathbf{Expr}\} \cup \{\mathbf{skip}\}$ be the set of *base statements*. We use $stmt$ to range over base statements.

Formally, a *parallel flow graph* comprises a finite set \mathbf{Proc} of *procedure names* that contains a distinguished procedure $Main$. Intuitively, $Main$ is the procedure with which execution starts. For simplicity, we assume that all procedures work on the same set X of global program variables and do not have local variables. Each procedure name $p \in \mathbf{Proc}$ is associated with a control flow graph $G_p = (N_p, E_p, A_p, e_p, r_p)$ that consists of:

- a set N_p of *program points*;
- a set of edges $E_p \subseteq N_p \times N_p$;
- a mapping $A_p : E_p \rightarrow \mathbf{Stmt} \cup \mathbf{Proc} \cup \mathbf{Proc}^2$ that annotates each edge with a base statement, a call of a single procedure, or a parallel call of two procedures; and
- a special *entry (or start) point* $e_p \in N_p$ and a special *return point* $r_p \in N_p$.

We assume that the program points of different procedures are disjoint: $N_p \cap N_q = \emptyset$ for $p \neq q$. This can always be enforced by renaming program points.

We write N for $\bigcup_{p \in \mathbf{Proc}} N_p$, E for $\bigcup_{p \in \mathbf{Proc}} E_p$, and A for $\bigcup_{p \in \mathbf{Proc}} A_p$. We also agree that $\mathbf{Base} = \{e \mid A(e) \in \mathbf{Stmt}\}$ is the set of base edges, $\mathbf{Call}_p = \{e \mid A(e) = p\}$ is the set of edges that call procedure p , and $\mathbf{Pcall}_{p,q} = \{e \mid A(e) = (p, q)\}$ is the set of edges that call procedure p and q in parallel. Moreover, we write \mathbf{Call} for $\bigcup_{p \in \mathbf{Proc}} \mathbf{Call}_p$ and \mathbf{Pcall} for $\bigcup_{p,q \in \mathbf{Proc}} \mathbf{Pcall}_{p,q}$.




Fig. 1. An example of a parallel flow graph.

Example 1 Figure 1 shows an example parallel flow graph with three procedures, *Main*, *p*, and *q*. The entry state of each procedure is marked by an arrow and the return state is indicated by a doubly circled state. The edge annotation **skip** is suppressed for clarity.

The main procedure of the example flow graph sequentially starts procedures *p* and *q*. Procedure *p* sets variable *y* to an arbitrary non-negative value and initializes *x* by 0. Procedure *q* has a choice: it can execute either the upper path, where it starts two new instances of *q* in parallel or the lower path, where it increments *x* by 2. Note that arbitrarily many instances of *q* can run in parallel. Upon termination *y* can hold an arbitrary non-negative number and *x* can hold an arbitrary non-negative number that is even. \square


The purpose of the remainder of this section is to set up a number of constraint systems, the solutions of which capture certain sets of program executions. In the next section we define an operational semantics that is useful as a reference point for setting up these constraint systems correctly.

2.2 Operational Semantics


We define a symbolic operational semantics of parallel flow graphs that specifies possible sequences of atomic actions. The evaluation of base statements is not described in this semantics. Thus, the configurations of the operational semantic represent control information only. In a sequential flow graph control information is simply given by a single flow-graph node. In a sequential program with procedures configurations would consist of sequences of flow-graph nodes. Such a sequence would model a stack of return addresses (or rather return nodes). In parallel flow graphs procedures can also be called in parallel. We model this by generalizing configurations from sequences to trees. Each node of the tree is labeled by a flow-graph node. Each inner node of the tree has either degree one—such nodes correspond to return addresses from simple calls or to return addresses from parallel calls where one of the parallel threads has terminated already—or degree two—such nodes correspond to return addresses from parallel calls. The active control points are given by

the leaves of the tree. Correspondingly, transitions are induced by the leaves. Transitions are labeled by base edges e , procedure names p , pairs of procedure names $p_0\|p_1$, or the symbol **ret**. There are four transition rules:


Base Step Rule: $c \xrightarrow{e} c'$, if $e = (u, v) \in \mathbf{Base}$ and c' results from c by replacing a leaf labeled u by a leaf labeled v .




Simple Call Rule: $c \xrightarrow{p} c'$, if there is an edge $e = (u, v) \in \mathbf{Call}_p$ such that c' results from c by replacing a leaf labeled u by a tree consisting of two nodes, a root labeled v and a successor node of the root labeled e_p .



Parallel Call Rule: $c \xrightarrow{p_0\|p_1} c'$, if there is an edge $e = (u, v) \in \mathbf{Call}_{p_0, p_1}$ such that c' results from c by replacing a leaf labeled u by a tree consisting of three nodes, a root labeled v with two successor nodes labeled e_{p_0} and e_{p_1} .



Return Rule: $c \xrightarrow{\mathbf{ret}} c'$, if c' results from c by removing a leaf labeled by r_p for some $p \in \mathbf{Proc}$.



When the Return Rule is applied the father of the node labeled r_p may become a leaf and thus become active. This models a return to a stacked return address. Just as well, however, the father may still have a child if it has degree two in c as indicated by the dotted line in the picture. In this case it becomes active only after the second leaf also vanishes. This models synchronized termination of threads started by a parallel call.

Note that the application of the Return Rule to a tree consisting of just a root results in the empty tree. Such a step models overall termination.

Let \mathbf{Conf} be the set of configurations, i.e., trees the degree of which is bounded by two and in which each node is annotated by a program point $u \in N$. We identify each program point $u \in N$ with the tree consisting of just a root labeled with u . We also write \mathbf{nil} for the empty tree. A program point $u \in N$ is *active* in a configuration c , if it labels one of the leaves of c . The predicate $At_u(c)$ is true if u is active in c and false otherwise.

Let $\mathbf{Label} = \mathbf{Base} \cup \mathbf{Proc} \cup \mathbf{Proc}^2 \cup \{\mathbf{ret}\}$ be the set of transition labels and $\longrightarrow \subseteq \mathbf{Conf} \times \mathbf{Label} \times \mathbf{Conf}$ be the transition relation defined by the rules above. We define the transitive generalization $\Longrightarrow \subseteq \mathbf{Conf} \times \mathbf{Label}^* \times \mathbf{Conf}$ of \longrightarrow , by

$$\xrightarrow{\varepsilon} = \mathbf{Id} \quad \xrightarrow{r \cdot \langle l \rangle} = \xrightarrow{r}; \xrightarrow{l},$$

where ‘;’ denotes relational composition, and write \xRightarrow{r} for $\bigcup_{r \in \mathbf{Label}^*} \xrightarrow{r}$. Here and in the following we write ε for the empty sequence, $\langle e_1, \dots, e_k \rangle$ for the sequence of the elements e_1, \dots, e_k , and \cdot for the concatenation operator.

2.3 Atomic Runs

As procedures do not have local variables, only the base edge labels in a transition sequence are of interest for dependence analysis. The other labels (calls, parallel calls, and returns) that appear between these labels can be ignored without losing interesting information. Therefore, we can abstract transition sequences to sequences of base edges safely. We call a sequence of base edges an (*atomic*) *run*; the set of atomic runs is $\mathbf{Runs} = \mathbf{Base}^*$. The classification ‘atomic’ refers to the fact that flow graph edges constitute atomic entities of execution; in Section 3 we consider ‘non-atomic runs’. We define for a label sequence l , \hat{l} to be the run obtained from l by retaining just the base edges and removing everything else:

$$\hat{\varepsilon} = \varepsilon \quad \text{and} \quad \widehat{r \cdot \langle l \rangle} = \begin{cases} \hat{r} \cdot \langle l \rangle & \text{if } l \in \mathbf{Base} \\ \hat{r} & \text{otherwise} \end{cases} \quad \text{for } r \in \mathbf{Label}^*, l \in \mathbf{Label}.$$

In the following we are going to set up constraint systems for a variety of run sets. These constraint systems use the following small number of operators and constants on run sets.

Semantics of base edges: $\llbracket e \rrbracket = \{\langle e \rangle\}$ for $e \in \mathbf{Base}$. This characterizes the run induced by a base edge in isolation.

Sequential composition operator: $R; S = \{r \cdot s \mid r \in R, s \in S\}$. This characterizes the sequential composition of run sets.

Interleaving operator: In order to define the interleaving (or parallel composition) operator some notation is needed. Let $r = \langle e_1, \dots, e_n \rangle$ be a sequence and $I = \{i_1, \dots, i_k\}$ a subset of positions in r such that $1 \leq i_1 < i_2 < \dots < i_k \leq n$. Then $r|I$ is the sequence $\langle e_{i_1}, \dots, e_{i_k} \rangle$. We write $|r|$ for the length of r , viz. n .

Then the interleaving of R and S is defined by

$$R \otimes S = \{r \mid \exists I_R, I_S : I_R \cup I_S = \{1, \dots, |r|\}, I_R \cap I_S = \emptyset, \\ r|I_R \in R, r|I_S \in S\}.$$

Prefix operator: $pre(R) = \{r \mid \exists s : r \cdot s \in R\}$. This captures prefixes of the runs in R .

Postfix operator: $post(R) = \{r \mid \exists s : s \cdot r \in R\}$. This captures postfixes of the runs in R .

Atomic runs may also be defined as sequences of base statements instead of base edges. For this we only need to redefine **Runs** as **Stmt*** instead of **Base*** and $\llbracket e \rrbracket$ by $\llbracket e \rrbracket = \{\langle A(e) \rangle\}$. In this setting we should also redefine the hat-operator to incorporate the transition from base edges to base statements:

$$\hat{\varepsilon} = \varepsilon \quad \text{and} \quad \widehat{r \cdot \langle l \rangle} = \begin{cases} \hat{r} \cdot \langle A(l) \rangle & \text{if } l \in \text{Base} \\ \hat{r} & \text{otherwise} \end{cases} \quad \text{for } r \in \text{Label}^*, l \in \text{Label}.$$

The remainder of this section can be read with both interpretations.

Non-standard semantics can be obtained by redefining the above operators. This is used in Section 3 for defining a semantics for parallel flow graphs in which execution of base edges is no longer assumed to be atomic. If we redefine the operators on an abstract domain with a finite chain height, we can effectively solve the constraint systems to be introduced soon by fixpoint iteration. If all these operators are correct or even precise abstractions of the concrete operators on atomic or non-atomic run sets, standard abstraction theorems from abstract interpretation ensure that the solution we get is a correct or even precise abstraction of the run sets characterized by the constraint systems. This is the idea of constraint-based program analysis.

2.4 The Run Sets of Ultimate Interest

We are ultimately interested in setting up constraint systems that characterize for each $u \in N$ the following sets of runs:

Reaching runs: $R(u) = \{\hat{r} \mid e_{Main} \xrightarrow{r} c, At_u(c)\}$.

Terminating runs: $T(u) = \{\hat{r} \mid e_{Main} \Longrightarrow c \xrightarrow{r} \text{nil}, At_u(c)\}$.

In dataflow analysis one considers *forward*- and *backward*-analyses. Forward-analyses calculate abstractions of the reaching runs and backward-analyses abstractions of the terminating runs.

We are also interested for all program points $u, v \in N$ in the set of those runs that potentially transfer information from u to v . We call these the *bridging runs from u to v* .

Bridging runs: $B_v(u) = \{\hat{r} \mid e_{Main} \Longrightarrow c_u \xrightarrow{r} c_v, At_u(c_u), At_v(c_v)\}$.

In the sections that follow, we present constraint systems that characterize the above run sets. That is: the smallest solution of these constraint systems consists of the run sets defined above. In addition to the above run sets, auxiliary run sets are necessary in order to formulate these constraint systems. These auxiliary run sets are stepwise introduced. We always explain the underlying intuition and outline the correctness proof but leave the details of the proof to the reader. The constraint systems for same-level, reaching and terminating runs are essentially taken from [4] where, however, they are not justified with reference to an explicitly given underlying operational semantics. The constraint system for bridging runs is new.

2.5 The Constraint Systems

2.5.1 Same-Level Runs

First of all, we characterize so-called *same-level runs*. Same-level runs of procedures capture complete runs of procedures in isolation.

Same-level runs of procedures: $S(q) = \{\hat{r} \mid e_q \xrightarrow{r} \text{nil}\}$ for $q \in \text{Proc}$.

As auxiliary sets we consider same-level runs to program nodes.

Same-level runs to program nodes: $S(u) = \{\hat{r} \mid e_q \xrightarrow{r} u\}$ for $u \in N_q$, $q \in \text{Proc}$.

Same-level runs of procedures form an important building block for the other constraint systems. Note that the complete effect of a parallel call edge $e \in \text{Pcall}_{p_0, p_1}$ is obtained easily from the same-level runs of procedures p_0 and p_1 : it is given by $S(p_0) \otimes S(p_1)$.

The same-level runs of procedures and program nodes are the smallest solution of the following constraint system:

$$\begin{aligned}
[\text{S1}] \quad & S(q) \supseteq S(r_q) \\
[\text{S2}] \quad & S(e_q) \supseteq \{\varepsilon\} \\
[\text{S3}] \quad & S(v) \supseteq S(u); \llbracket e \rrbracket, \quad \text{if } e = (u, v) \in \text{Base} \\
[\text{S4}] \quad & S(v) \supseteq S(u); S(p), \quad \text{if } e = (u, v) \in \text{Call}_p \\
[\text{S5}] \quad & S(v) \supseteq S(u); [S(p_0) \otimes S(p_1)], \text{ if } e = (u, v) \in \text{Pcall}_{p_0, p_1}
\end{aligned}$$

It is easy to see that the same-level runs satisfy all constraints:

- [S1]: A same-level run of the return point of procedure q gives rise to a same-level run of q by the Return Rule.
- [S2]: It follows trivially from the definition that ε is a same-level run of the entry point of a procedure.
- [S3]: If $e = (u, v)$ is a base edge, we get a same-level run to v by extending a same-level run to u with e by the Base Steps Rule.
- [S4]: If $e = (u, v)$ is an edge that calls p , we get a same-level run to v if we extend a same-level run to u by a same-level run of p : we follow the execution underlying the same-level run to v and then call p according to the Simple Call Rule; we then follow the execution underlying the same-level run of p (with v waiting on the stack to become active) and return to v according to the Return Rule.
- [S5]: Similarly, if $e = (u, v)$ is an edge that calls p_0 and p_1 in parallel, we can—after seeing a same-level run to u —follow this edge; then p_0 and p_1 are performed to completion in parallel, which results in an interleaving of a same-level run of p_0 and p_1 ; after that, execution returns to v . We thus obtain a same-level run to v by extending a same-level run of u with an interleaving of same-level runs of p_0 and p_1 .

On the other hand, we can easily prove by induction on the length of the transition sequences inducing same-level runs, that each same-level run lies in any solution of the constraint system, in particular in the smallest one: in the base case we consider the empty execution ε . It can only give rise to the same-level run ε to e_q for some procedure q . But ε is enforced to lie in any solution of $S(r_p)$ explicitly by constraint [S2].

In the induction step, we consider longer executions leading to same-level runs. The execution underlying a same-level run of a procedure q necessarily involves a final return from r_q after an execution that gives rise to a same-level run of r_q . The latter execution is one step shorter and thus the same-level run of r_q is contained in any solution of $S(r_q)$ by the induction hypothesis. Now, the constraint [S1] ensures that it is also contained in the set assigned to $S(q)$ in a solution.

The last step of a non-empty execution r inducing a same-level run \hat{r} to a program point v must be induced either by the Base Rule or the Return Rule because the Simple and Parallel Call Rule never lead to a configuration which consists of just a single state. If the last step is induced by the Base Rule, the previous configuration is a program point u . Then \hat{r} is composed of a same-level run to u and the base edge $e = (u, v)$. The same-level run to u is induced by a shorter execution and hence contained in the set associated with $S(u)$ in any solution by the induction hypothesis. Thus, \hat{r} is in $S(v)$ by the constraint [S3]. If the last step is induced by the Return Rule, then there must be a simple or parallel call from which this step returns. The constraints for simple and parallel call edges ([S4] and [S5]) together with the induction hypothesis then ensure that \hat{r} is contained in $S(v)$.

2.5.2 Inverse Same-Level Runs

We also consider a kind of dual to same-level runs of program points: runs from a program point to the return point of the corresponding procedure. We call these *inverse same-level runs of program point*. They are needed in order to capture terminating runs.

Inverse same-level runs of program points:

$$S^i(u) = \{\hat{r} \mid u \xrightarrow{\hat{r}} \text{nil}\} \text{ for } u \in N.$$

Inverse same-level runs of procedures and program nodes are obtained by backwards accumulation as the smallest solution of the following system of constraints:

$$\begin{aligned} \text{[SI1]} \quad & S^i(r_q) \supseteq \{\varepsilon\} \\ \text{[SI2]} \quad & S^i(u) \supseteq \llbracket e \rrbracket; S^i(v), \quad \text{if } e = (u, v) \in \text{Base} \\ \text{[SI3]} \quad & S^i(u) \supseteq S(p); S^i(v), \quad \text{if } e = (u, v) \in \text{Call}_p \\ \text{[SI4]} \quad & S^i(u) \supseteq [S(p_0) \otimes S(p_1)]; S^i(v), \text{ if } e = (u, v) \in \text{Pcall}_{p_0, p_1} \end{aligned}$$

The last two constraints refer to same-level runs of procedures. Therefore, it appears that we need to calculate same-level runs before we can calculate inverse same-level runs by the above constraint system. However, by adding for each procedure $q \in \text{Proc}$ the constraint

$$\text{[SI5]} \quad S(q) \supseteq S^i(e_q)$$

we can calculate same-level runs of procedures simultaneously with inverse same-level runs. Thus, we can also calculate inverse same-level runs in isolation.

It is easy to see that the sets of inverse same-level runs satisfy all constraints:

- [SI1]: By the Return rule, ε clearly is an inverse same-level run of the return point r_q of a procedure.
- [SI2]: If $e = (u, v)$ is a base edge, we get an inverse same-level run of u by prefixing a same-level run of v with e .
- [SI3]: If $e = (u, v)$ is an edge that calls p , we can follow this edge in an execution from u ; then p is performed until termination, which results in a same-level run of p ; after that execution proceeds at v . We thus obtain an inverse same-level run of u by prefixing an inverse same-level run of v by a same-level run of p .
- [SI4]: Similarly, if $e = (u, v)$ is an edge that calls p_0 and p_1 in parallel, we can follow this edge in an execution from u ; then p_0 and p_1 are performed to completion in parallel, which results in an interleaving of a same-level run of p_0 and p_1 ; after that execution returns to v . We thus obtain an inverse same-level run of u by prefixing an inverse same-level run of v with an interleaving of same-level runs of p_0 and p_1 .

On the other hand, we can easily prove by induction on the length of the transition sequences inducing inverse same-level runs, i.e. those that lead to nil, that each inverse same-level run is in the smallest solution of the constraint system: in the base case we consider the shortest executions that lead to same-level runs. These are executions of the form $r_p \xrightarrow{\text{ret}} \text{nil}$ for some procedure p . They witness that $\varepsilon \in \mathbf{S}(r_p)$. But ε is enforced to be in a solution of $S(r_p)$ explicitly by constraint [SI1].

In the induction step, we consider longer executions leading to same-level runs. These necessarily start with a transition induced by a base edge, a simple, or a parallel call edge. The resulting run is then composed from shorter runs as specified in the constraints for base edges ([SI2]), simple calls ([SI3]), and parallel calls ([SI4]), respectively.

2.5.3 Two Assumptions and a Simple Analysis

The following two assumptions simplify the constraint systems that follow:

ASS1: every program point $u \in N_q$ in a procedure q can be reached by a same-level run from the entry point e_q of q :


$$\forall q \in \text{Proc}, u \in N_q : \mathbf{S}(u) \neq \emptyset.$$

ASS2: from every program point $u \in N_q$ the return point r_q can be reached by a same-level run:

$$\forall q \in \text{Proc}, u \in N_q : \mathbf{S}^i(u) \neq \emptyset.$$

These assumptions are not as innocent as they may seem at first glance. In particular it does *not* suffice to require that there are paths from e_q to u and

from u to r_q in the flow graph G_q for q . The paradigmatic counter-example is a procedure that calls itself and has no bypassing terminating branch:



Although there is a path from e_q to r_q in the flow graph, no execution can reach r_q from e_q , as there is no terminating bypass of the recursive call of q . Hence both $\mathbf{S}(r_q)$ and $\mathbf{S}^i(e_q)$ are empty. Examples like this show that we cannot assume without loss of generality that practical flow graphs satisfy ASS1 and ASS2.

While assumptions ASS1 and ASS2 simplify the presentation and justification of the constraint systems in the remainder of this section, they are not strictly necessary. We explain the necessary changes for the general case in Section 2.5.7.

In order to compute the information needed to decide ASS1 and ASS2, we design a simple analysis procedure. It is based on an abstract interpretation of the operators and constants used in the constraint systems. We work with a two point domain ($\mathbb{D} = \{\perp, \top\}, \leq$) ordered as $\perp \leq \top$. The idea is that \perp represents definite emptiness and \top potential non-emptiness of a run set. Correspondingly, we define the abstraction mapping $\alpha : 2^{\text{Runs}} \rightarrow \mathbb{D}$ by $\alpha(\emptyset) = \perp$ and $\alpha(R) = \top$ for $R \neq \emptyset$. The fact that the abstract interpretation developed below is precise guarantees that it computes indeed \perp for all empty run sets and \top just for non-empty run sets. Obviously, α is universally disjunctive. We define the abstract interpretation of the operators by

$$x;^\#y = x \otimes^\# y = x \wedge y, \quad \text{pre}^\#(x) = \text{post}^\#(x) = x, \quad \llbracket e \rrbracket^\# = \{\varepsilon\}^\# = \top$$

for $x, y \in \mathbb{D}$, $e \in E$. It is easy to see that the abstract operators are precise abstractions of the corresponding operators on run sets: a sequential or parallel composition of two run sets is non-empty iff both arguments are non-empty; the set of prefixes and the set of postfixes of a run set R are non-empty iff R is; and each base edge gives rise to a non-empty run set. Therefore, by computing the least solution of the constraint systems for same-level and inverse same-level runs over the abstract interpretation we get precise information about the emptiness of the sets of same-level and inverse same-level runs of program points.

This analysis is cheap: as (\mathbb{D}, \leq) has chain height two, the information for each constraint variable can change at most once in the fixpoint iteration. By standard demand-driven fixpoint evaluation, we can organize the computation of the least solution such that each operator in the constraint system is evaluated at most once. Thus, the computation can be done in time $\mathcal{O}(|E| + |\text{Proc}|)$, the

number of operators in the constraint systems. As in all practical flow graphs out-degrees of program nodes are bounded, typically by 2, and $|\text{Proc}|$ is trivially bounded by $|N|$ as each procedure has a distinguished entry node, this typically is $\mathcal{O}(|N|)$. In the following we assume that this analysis has been done such that for each program node u and procedure q the information whether $S(u)$, $S^i(u)$, $S(q)$, or $S^i(q)$ is empty or not is readily available.

Another analysis that can determine information about reachability of program points in parallel flow graphs has been described by Seidl and Steffen [4] as an instance of their generic analysis framework for solving gen/kill dataflow problems for parallel programs.

2.5.4 Reaching Runs

As auxiliary sets for characterizing the runs that reach a program point u , we consider the runs that reach u from a call to procedure q .

Reaching runs from procedures: $R(u, q) = \{\hat{r} \mid e_q \xrightarrow{r} c, At_u(c)\}$ for $u \in N, q \in \text{Proc}$.

With this definition, we obviously have $R(u) = R(u, \text{Main})$. Hence we are done with characterizing reaching runs if we succeed in characterizing reaching runs from procedures. The latter can be done by the following constraint system:

$$\begin{aligned} [\text{R1}] \quad R(u, q) &\supseteq S(u), && \text{if } u \in N_q \\ [\text{R2}] \quad R(u, q) &\supseteq S(v); R(u, p), && \text{if } (v, _) \in E_q \cap \text{Call}_p \\ [\text{R3}] \quad R(u, q) &\supseteq S(v); [R(u, p_i) \otimes \text{pre}(S(p_{1-i}))], && \text{if } (v, _) \in E_q \cap \text{Pcall}_{p_0, p_1} \end{aligned}$$


The last clause is meant to specify two constraint for $i = 0$ and $i = 1$.

The reaching runs satisfy the constraints:

- [**R1**]: Firstly, each same-level run of u clearly is also a reaching run of u .
- [**R2**]: Secondly, if we have a program point v in q that has an outgoing edge calling p —the situation described in the second constraint—we obtain a run that reaches u from q when we extend a same-level run \hat{r} to v with a run \hat{r}' that reaches u from p (where r and r' are the underlying executions).
- [**R3**]: Thirdly, consider a program point v in q that has an outgoing edge calling p_0 and p_1 in parallel, the situation described in the third constraint. Similar to the second case, we get a run reaching u by extending a same-level run of v with a run that reaches u in the parallel call. The latter can happen either in p_0 or p_1 hence the two cases with $i = 0, 1$. Now until p_i has reached u in p_i the other procedure p_{1-i} can perform a prefix of a same-level run.

On the other hand, the constraint system captures all the ways how u may be reached from e_q . There are just three possibilities: either u is on the same-level, in a simple call, or in a parallel call. These case are completely covered by the constraints.

Note that assumption ASS2 is crucial for making the constraint for parallel calls sufficiently rich. If it is violated, the partial run exhibited by p_{1-i} while p_i is in the process of reaching u need not be a prefix of a same-level run. For example, the following procedure q might execute $x := e$ arbitrarily often, although $S(q)$ and hence $pre(S(q))$ is empty.



A possible remedy is described in Section 2.5.7.

2.5.5 Terminating Runs

The approach for capturing terminating runs is dual to the one for reaching runs. As auxiliary sets we consider terminating runs of u in a call to procedure q .

Terminating runs in procedures: $\mathbb{T}(u, q) = \{\hat{r} \mid e_q \Longrightarrow c \xrightarrow{r} \text{nil}, At_u(c)\}$
for $u \in N, q \in \text{Proc}$.

Obviously we have $\mathbb{T}(u) = \mathbb{T}(u, \text{Main})$ such that it suffices to capture terminating runs in procedures in the constraint system. The constraint system is dual to the one for reaching runs:

- [T1] $T(u, q) \supseteq S^i(u)$, if $u \in N_q$
- [T2] $T(u, q) \supseteq T(u, p); S^i(w)$, if $(-, w) \in E_q \cap \text{Call}_p$
- [T3] $T(u, q) \supseteq [T(u, p_i) \otimes post(S(p_{1-i}))]; S^i(w)$, if $(-, w) \in E_q \cap \text{Pcall}_{p_0, p_1}$

Again, $i = 0, 1$ in the last constraint. The justification of this constraint system is similar to reaching runs; therefore, the details are left to the reader. We should mention, however, that assumption ASS1 is crucial here, like ASS2 in the case of reaching runs, but for a quite different reason. The difference is the requirement that the configuration c with $At_u(c)$ is reachable ($e_q \Longrightarrow c$) in terminating runs, a requirement that has no analogue for reaching runs. As a consequence, $post(S(p_{1-i}))$ is now sufficient to capture the interleaving potential in the constraint for parallel calls even in the general case, in contrast to $pre(R(p_{1-i}))$ in the corresponding constraint for reaching runs.

However, the reachability requirement for configuration c , implies that some of the constraints are not satisfied by the sets $\mathbb{T}(u, q)$ in the general case. For example, an inverse same-level run r from a program point $u \in N_q$ is not always a terminating run. Being an inverse same-level run just means that $u \xrightarrow{r} \text{nil}$ holds, but for a terminating run we additionally need $e_q \Longrightarrow u$. This is automatically true if ASS1 is valid but can be wrong in the general case. Similarly, we need that the start node of the edge e in the second and third constraint can be reached for making the constraints valid for the operationally defined sets. A possible remedy is to remove the constraints induced by non-reachable program points. This is detailed in Section 2.5.7.

2.5.6 Bridging Runs

Let $v \in N$ be a fixed program point. We want to determine the bridging runs $B_v(u)$ for each $u \in N$ as defined in Section 2.4. As a first step we capture for each program points u the runs that reach v , when execution is started directly with u . We call these the *simple bridging runs* of u w.r.t. v .

Simple bridging runs: $B_v^s(u) = \{\hat{r} \mid u \xrightarrow{r} c, At_v(c)\}$ for $u \in N$.

The simple bridging runs can be characterized as the smallest solution of the following constraint system:

$$\begin{aligned}
[\text{BS1}] \quad & B_v^s(v) \supseteq \{\varepsilon\} \\
[\text{BS2}] \quad & B_v^s(u) \supseteq \llbracket e \rrbracket ; B_v^s(w), & \text{if } e = (u, w) \in \text{Base} \\
[\text{BS3}] \quad & B_v^s(u) \supseteq S(p) ; B_v^s(w), & \text{if } e = (u, w) \in \text{Call}_p \\
[\text{BS4}] \quad & B_v^s(u) \supseteq B_v^s(e_p), & \text{if } e = (u, _) \in \text{Call}_p \\
[\text{BS5}] \quad & B_v^s(u) \supseteq [S(p_0) \otimes S(p_1)] ; B_v^s(w), & \text{if } e = (u, w) \in \text{Pcall}_{p_0, p_1} \\
[\text{BS6}] \quad & B_v^s(u) \supseteq B_v^s(e_{p_i}) \otimes \text{pre}(S(p_{1-i})), & \text{if } e = (u, _) \in \text{Pcall}_{p_0, p_1}
\end{aligned}$$

The last constraint is again included for $i = 0, 1$.

Let us explain why these constraints cover all the ways how v can be reached from u . If $u = v$ then there is the trivial way to reach v from u : by the empty execution; this is covered by Constraint [BS1]. Otherwise, we must proceed via an outgoing edge (u, w) of u . If this is a base edge $e = (u, w)$, we first see e and then a run that reaches v from w ; this is covered by Constraint [BS2]. If e is an edge that calls a procedure p , we distinguish two cases: either v is reached after p has terminated—this case is covered by Constraint [BS3]—or v is reached during the execution of p —this case is covered by [BS4]. Similarly, if e is a parallel call of two procedures p_0 and p_1 , we can reach v either after both procedures have terminated, which is covered by [BS5]. Or we can reach v in one of the called procedures p_i . In this case we see a run from e_{p_i} that reaches v interleaved with a prefix of a same-level run of procedure p_{1-i} . If assumption

ASS2 is violated we must again reckon with procedure p_{1-i} providing runs that are not prefixes of same-level runs, as was the case for reaching runs. We can solve this problem as for reaching runs, cf. Section 2.5.7.

The reader should face no difficulties in persuading himself, that the $B_v^s(u)$ sets indeed solve all constraints.

As a second step we determine the bridging runs in a call to a procedure:

Bridging runs in procedure calls:

$$B_v(u, q) = \{\hat{r} \mid e_q \Longrightarrow c_u \xrightarrow{r} c_v, At_u(c), At_v(c)\} \text{ for } u \in N.$$

Clearly, we have $B_v(u) = B_v(u, Main)$ such that we are done, when we have successfully captured $B_v(u, q)$ for all u, q .

Basically, there are two ways how a bridging run may occur in a call to q . One possibility is that both u and v are reached in the same simple or parallel call in q . This case is captured by the following three types of constraints:

$$\begin{aligned} \text{[B1]} \quad B_v(u, q) &\supseteq B_v(u, p), && \text{if } e \in E_q \cap \text{Call}_p \\ \text{[B2]} \quad B_v(u, q) &\supseteq B_v(u, p_i) \otimes \text{post}(\text{pre}(S(p_{1-i}))), && \text{if } e \in E_q \cap \text{Pcall}_{p_0, p_1} \\ \text{[B3]} \quad B_v(u, q) &\supseteq \text{pre}(T(u, p_i)) \otimes \text{post}(R(v, p_{1-i})), && \text{if } e \in E_q \cap \text{Pcall}_{p_0, p_1} \end{aligned}$$

[B2] and [B3] apply for $i = 0, 1$.

Constraint [B1] captures the case that u and v are reached in the same simple call. Constraint [B2] is concerned with the case that u and v are reached in the same procedure p_i of a parallel call. Before u is reached in p_i the other procedure can already perform certain actions and it need not run to completion until v is reached. Therefore, p_{1-i} contributes a middle piece of a same-level run. Potential middle pieces can be characterized by $\text{pre}(\text{post}(S(p_{1-i})))$ as captured by the second constraint. Constraint [B3] captures the case that u is reached in procedure p_i and v in procedure p_{1-i} . After p_i has reached u it can further proceed; specifically p_i contributes a prefix of a run from $T(u)$ until v is reached in p_{1-i} . In order to reach v , p_{1-i} must execute a run from $R(v, p_{1-i})$. It can execute a prefix of this run before p_i leaves u . Therefore, we see a postfix of a run from $R(v, p_{1-i})$ as part of the bridging run.

The second possibility is that u and v are not reached in the same simple or parallel call. This gives rise to the following constraints:

$$\begin{aligned} \text{[B4]} \quad B_v(u, q) &\supseteq B_v^s(u), && \text{if } u \in N_q \\ \text{[B5]} \quad B_v(u, q) &\supseteq T(u, p); B_v^s(w), && \text{if } (-, w) \in E_q \cap \text{Call}_p \\ \text{[B6]} \quad B_v(u, q) &\supseteq [T(u, p_i) \otimes \text{post}(S(p_{1-i}))]; B_v^s(w), && \text{if } (-, w) \in E_q \cap \text{Pcall}_{p_0, p_1} \end{aligned}$$

where $i = 0, 1$ in the last constraint.

The first subcase is that u is reached on same-level, i.e. in the current instance of q . Then we see a simple bridging run of u (Constraint [B4]). The second subcase is that u is reached in a procedure p called by a simple call edge $e = (-, v) \in E_q$. Then we see a run from $T(u, p)$ followed by a simple bridging run from w (Constraint [B5]). The third subcase is that u is reached in a procedure p_i called by a parallel call edge $e = (-, v) \in E_q$. Then we see a run from $T(u, p_i) \otimes \text{post}(S(p_{1-i}))$ followed by a simple bridging run from w (Constraint [B6]).

2.5.7 The General Case

In this section we describe the changes that are necessary in the general case, i.e., if assumptions ASS1 and ASS2 are potentially violated.

As explained in connection with constraint [R3] one of the problems is that in the general case $\text{pre}(S(q))$ does not capture all partial runs of procedure q . Thus, interleaving $R(u, p_i)$ with $\text{pre}(S(p_{1-i}))$ does not capture all possible run that reach u in a parallel call. This problem also arises in constraints [BS6] and [B2]. A possible remedy is to introduce new variables $P(q)$, $q \in \text{Proc}$, that characterize finite prefixes of (finite or infinite) runs, i.e. $P(q) = \{\hat{r} \mid e_q \xrightarrow{r} c\}$, and to use $P(p_{1-i})$ instead of $\text{pre}(S(p_{1-i}))$ in [R3], [BS6], and [B2]. A simple way to calculate $P(q)$ is to add a constraint of the following form for each procedure q and program point u to the constraint system for reaching runs:³

$$[P] P(q) \supseteq \text{pre}(R(u, q)).$$

While this way of calculating $P(q)$ is easy to specify it has the disadvantage of introducing $|N| \cdot |\text{Proc}|$ new constraints, i.e. quadratically many. Although this does not spoil the overall asymptotic complexity—already the constraint system for reaching runs has $\mathcal{O}(|N| \cdot |\text{Proc}|)$ constraints—we should mention that $P(q)$ can be calculated also by $\mathcal{O}(|N|)$ constraints. A corresponding constraint system is given in Fig. 2. It determines as auxiliary information finite prefixes of (finite or infinite) runs from program points, defined by $P(u) = \{\hat{r} \mid u \xrightarrow{r} c\}$ by backwards accumulation and is similar to the constraint system for simple bridging runs.

A similar problem arises in constraint [B3]: if assumption ASS2 is violated,

³ For the atomic case the pre -operator may be omitted as any configuration c satisfies $\text{At}_u(c)$ for at least one program point u . In the non-atomic interpretation, however, there are (implicitly) *transient* configurations that correspond to intermediate stages of executions in which no program point is active. Fortunately, from all transient configurations c a configuration c' with some active program point is reachable. Therefore, we can capture the runs to transient configurations by means of the pre-operator.

$$\begin{aligned}
[\text{P1}] \quad & P(q) \supseteq P(e_q) \\
[\text{P2}] \quad & P(u) \supseteq \{\varepsilon\} \\
[\text{P3}] \quad & P(u) \supseteq \llbracket e \rrbracket; P(v), \quad \text{if } e = (u, v) \in \text{Base} \\
[\text{P4}] \quad & P(u) \supseteq P(p), \quad \text{if } (u, _) \in \text{Call}_p \\
[\text{P5}] \quad & P(u) \supseteq S(p); P(v), \quad \text{if } (u, v) \in \text{Call}_p \\
[\text{P6}] \quad & P(u) \supseteq [P(p_0) \otimes P(p_1)], \quad \text{if } (u, v) \in \text{Pcall}_{p_0, p_1} \\
[\text{P7}] \quad & P(u) \supseteq [S(p_0) \otimes S(p_1)]; P(v), \quad \text{if } (u, v) \in \text{Pcall}_{p_0, p_1}
\end{aligned}$$

Fig. 2. A constraint system characterizing finite prefixes.

$$\begin{aligned}
[\text{Q1}] \quad & Q(u, q) \supseteq P(u), \quad \text{if } u \in S_q \\
[\text{Q2}] \quad & Q(u, q) \supseteq Q(u, p), \quad \text{if } (v, _) \in E_q \cap \text{Call}_p \\
[\text{Q3}] \quad & Q(u, q) \supseteq T(u, p); P(w), \quad \text{if } (v, w) \in E_q \cap \text{Call}_p \\
[\text{Q4}] \quad & Q(u, q) \supseteq Q(u, p_i) \otimes \text{post}(P(p_{1-i})), \quad \text{if } (v, _) \in E_q \cap \text{Call}_{p_0, p_1} \\
[\text{Q5}] \quad & Q(u, q) \supseteq [T(u, p_i) \otimes \text{post}(S(p_{1-i}))]; P(w), \quad \text{if } (v, w) \in E_q \cap \text{Call}_{p_0, p_1}
\end{aligned}$$

Fig. 3. A constraint system for partial runs that can be exhibited in a procedure after a given program point has been reached. All constraints [Q1]-[Q5] are only for program points v with $S(v) \neq \emptyset$. In [Q4] and [Q5], $i = 0, 1$.

$\text{pre}(T(u, p_i))$ does not necessarily capture all partial runs exhibited by p_i after reaching u because u could be reached at a configuration from which termination is impossible. The information needed in place of $\text{pre}(T(u, p_i))$ is $Q(u, p_i)$ where $Q(u, q) = \{\hat{r} \mid e_q \Longrightarrow c \xrightarrow{r} c', \text{At}_u(c)\}$ for $u \in N$, $q \in \text{Proc}$. These sets can be characterized by the constraint system in Fig. 3

The above changes ensure that the run sets characterized by the constraint systems are sufficiently large. They are necessary to make flow analysis based on abstract interpretation of the constraint systems sound. The changes described now ensure that the run sets do not become too large. Thus, they are necessary to make analyses based on a precise abstract interpretation complete.

As explained in connection with terminating runs, constraints induced by unreachable program points are not satisfied by the run sets (defined from the operational semantics) that we intend to characterize. As these constraints pose unnecessary additional requirements they make the solutions larger than necessary. Fortunately, such constraints are also unnecessary for soundness and can simply be removed. Specifically, we must include the constraints [T1], [B1], and [B4] only for program points u with $S(u) \neq \emptyset$, and the constraints [T2], [T3], [B2], [B3], [B5], and [B6] only for edges $e = (v, w)$ with $S(v) \neq \emptyset$.

We have seen in Section 2.5.3, that we can determine this information with a very simple and cheap analysis.

With the changes described in this section we obtain constraint systems that are both sound and complete for the general case.

2.6 Discussion

In this section we have introduced parallel flow graphs. After that we defined a symbolic operational semantics. It works on configurations that take the form of a tree, the nodes of which are annotated by program points. Such a tree models a generalization of a run time stack that may branch to parallel stacks in addition to the common stack operations. We have described the transitions of the operational semantics by rules that work directly on configurations of this form.

Alternatively, we could have used the approach of Esparza, Knoop, and Podelski [5,6]. They map a parallel flow graph to a so-called PA-processes; PA is a process algebra which has both a sequential and a concurrent composition operator [24,25]. Execution of PA-processes in turn is described by a structured operational semantics (SOS) [26]. This enable them to apply results about model-checking of PA-processes to flow analysis. For our purposes the approach chosen here is sufficient and produces less notational overhead.

Based on the operational semantics we have defined some run sets of interest and developed constraint systems that characterize these run sets. The constraint systems for same-level runs and reaching runs are essentially the ones used by Seidl and Steffen [4]. Also the constraint systems for inverse same-level runs and terminating runs are indicated in their work. The constraint system for bridging runs, however, is new. Seidl and Steffen *postulate* their constraint systems, while we use an operational semantics as a reference point. While this might be considered a minor or even trivial difference, in our opinion an operational justification of the constraint systems increases our understanding of what exactly is specified by the constraint systems.

Many reasonable variants of the run sets in question may be considered. For example, one could define reaching runs by

$$R'(u) = \{\hat{r} \mid e_{Main} \xrightarrow{r} c \implies \varepsilon, At_u(c)\}.$$

This definition deviates from our previous definition in that it considers only configurations c from which termination is possible, i.e., it characterizes the runs that both reach u and can be completed to a terminating run. If assumption ASS2 is violated, the new definition gives rise to smaller run sets. Sim-

ilarly, many reasonable variants of the other run sets are conceivable and by techniques similar to the ones of Section 2.5.7 sound and complete constraint systems for these variants can be constructed. Operational specifications of the run sets in question allow to distinguish these variants much more clearly than implicit specifications by means of constraint systems.

Validating constraint systems with respect to an operational semantics has another advantage: it helps to uncover subtle bugs. In the absence of an operational semantics Seidl and Steffen, for instance, fail to notice that constraint [R3] in the constraint system for reaching runs is not rich enough to characterize all reaching runs in a parallel composition if assumption ASS2 does not hold. We detected this error while trying to justify the soundness of the constraint system. As a consequence their constraint system for reaching runs is unsound in the general case. To be fair, we should note that this does not affect the soundness of their analysis procedure that is not directly based on the constraint system for reaching runs. We should also say that they solve the problems that arise when assumption ASS1 is violated correctly. Here they validly propose to remove edges leaving unreachable program points before the analysis. This has essentially the same effect as the side conditions of the form $S(u) \neq \emptyset$ added to the various constraints in Section 2.5.7.

3 Non-Atomic Execution

The idealization that assignments execute atomically is quite common in the literature on program verification as well as in the theoretical literature on flow analysis of parallel programs. However, in a multi-processor environment where a number of concurrently executing processors share a common memory this assumption is hardly realistic. In such an environment two threads of control may well interfere while each of them is in the process of executing an assignment. The reason is that assignments are broken into smaller instructions before execution.

As a simple example, consider a program in which a shared variable x is incremented by two threads in parallel:

$$x := x + 1 \parallel x := x + 1.$$

Let us assume that x holds 0 initially. If assignments execute atomically, this program clearly will increment x twice and so terminate in a state in which variable x holds 2. However, in a multi-processor environment this program may well set x to 1. For example, the following execution may happen: first, one of the processors accesses the memory in order to get the value of x . While it is in the process of incrementing this value, but before it has written back

the result, the second processors may access the memory, too, in order to get the value of x . In such a run, both processors read the initial value 0 for x , both will increment just this value, and both will write back 1 for x . Consequently, the program will terminate in a state where x holds 1 instead of 2.

The moral of this discussion is that, in the real world of multi-processor execution, we cannot assume atomic execution of assignments. What we typically *may* safely assume, however, is that single reads of variables and single writes of variables are atomic, because the access to the memory is usually synchronized, e.g., through a common bus. We can develop an interleaving semantics for parallel programs that adequately models non-atomic execution of assignments by means of breaking assignments into more fine-grained atomic actions, an observation that is exploited in a moment.

This said, we should mention that there are indeed execution scenarios for concurrent programs that guarantee atomic execution of assignments. In particular in a time-shared multi-tasking environment, where concurrent execution of threads is simulated by a single processor that switches between execution of code pieces implementing the different threads, assuming atomic execution of assignments may be safe, if context switches happen only between assignments, but not in the process of executing the code implementing a single assignment. The built-in scheduler of the Transputer, for instance, performs context switches only after certain types of instructions that typically end execution of assignment code [27].⁴

In this section we provide parallel flow graphs with an interleaving semantics that models non-atomic execution of assignments adequately. For this purpose we define a domain \mathbf{NR} of sets of (non-atomic) runs and provide adequate definitions for the constants and operators used in the constraint systems in Section 2.5. Specifically, we provide

- an interpretation $\llbracket e \rrbracket \in \mathbf{NR}$ for the non-atomic runs of a base edge; and
- interpretations for the operators $;$, \otimes , *pre*, and *post* used in the constraint systems.

Solving the constraint systems from Section 2.5 over this new interpretation immediately gives us adequate definitions for the reaching, terminating, and bridging runs of a parallel flow graph when assignments execute non-

⁴ The Transputer designers chose this strategy in order to make context switches cheap and fast. In typical code, the contents of certain registers used for expression evaluation is no longer needed after such instructions. Therefore, these registers are not stored during context switches, which makes context switches fast. Actually, it is the compiler writer's task to ensure that the generated code does not rely on the registers keeping their contents after such instructions. Atomic execution of assignments in typical code is a neat side-effect of this design.

atomically.

3.1 Modeling Non-Atomic Execution by Virtual Variables

Suppose given a parallel flow graph and let X be the set of *program variables* which the statements of the flow graph refer to. In order to explain the meaning of non-atomic statements appropriately suppose furthermore given an infinite set V of *virtual (or internal) variables* disjoint from X . Intuitively, virtual variables are used to store intermediate results that are private to the threads. The parallel composition (or interleaving) operator defined later ensures that parallel threads do not interfere on virtual variables. We use the letters x, y to range over X , u, v to range over V , and the letters a, b to range over $X \cup V$.

For the purpose of the semantics, assignments are split into atomic operations. As an example consider an assignment statement $x := e(y_1, \dots, y_k)$ in the program where y_1, \dots, y_k refer to the occurrences of (program) variables in expression e . There are many sensible atomicity assumptions. For example, we could work with the rather pessimistic assumption that just reads and writes of variables are atomic and that variables appearing more than once in e are re-read for every occurrences. Then $x := e(y_1, \dots, y_k)$ is replaced by a sequence of assignments

$$\langle v_{\pi(1)} := y_{\pi(1)}, \dots, v_{\pi(k)} := y_{\pi(k)}, x := e(v_1, \dots, v_k) \rangle,$$

where v_1, \dots, v_k are arbitrary distinct virtual variables and π is a permutation of $\{1, \dots, k\}$. The idea is that the other threads can execute atomic operations between these assignments.

More coarse-granular atomicity assumptions can be captured in a similar way. E.g. if we assume that the evaluation of the right-hand-side expression is atomic then we would replace $x := e(y_1, \dots, y_k)$ by

$$\langle v := e(y_1, \dots, y_k), x := v \rangle.$$

The important point to notice is that—whatever the specific atomicity assumption may be—if we assume that the execution of all assignments is non-atomic, then all assignments in a run that refer to a *program variable* on the left hand side have only *virtual variables* on the right hand side. Thus, all assignments belong to the set

$$\text{Asg} = \{a := e(b_1, \dots, b_k) \mid a \in X \Rightarrow b_1, \dots, b_k \in V\}.$$

One way of obtaining a semantics for non-atomically executing assignments is to transform the assignments in the program prior to semantic interpretation.




Fig. 4. Introduction of virtual variables.




Fig. 5. Confusion of virtual variables.

As an example consider the program in Fig. 4(a) which corresponds to the example discussed in the introduction. We could transform it to the program in Fig. 4(b) and then apply the standard interpretation.

The problem with this approach is that we must be careful not to confuse virtual variables of different threads. This is simple if only instances of different procedures run in parallel: then we can simply use different names for the virtual variables in different procedures. However, it becomes problematic if different instances of the same procedure may run in parallel like in the program in Fig. 5. Then we must model the virtual variables by local variables of the procedures which is not supported by the flow graph model developed up to now. Therefore, we use a different approach: instead of transforming flow graphs we incorporate the transformation implicitly into the semantic interpretation of assignments.

Before we turn to the technical details of the new semantic interpretation we present an example where the answer to the constant detection problem depends on the atomicity assumption for base statements. This example illustrates that the main mechanism underlying the undecidability proof of interprocedural parallel constant detection [12] does not carry over to the non-atomic case.

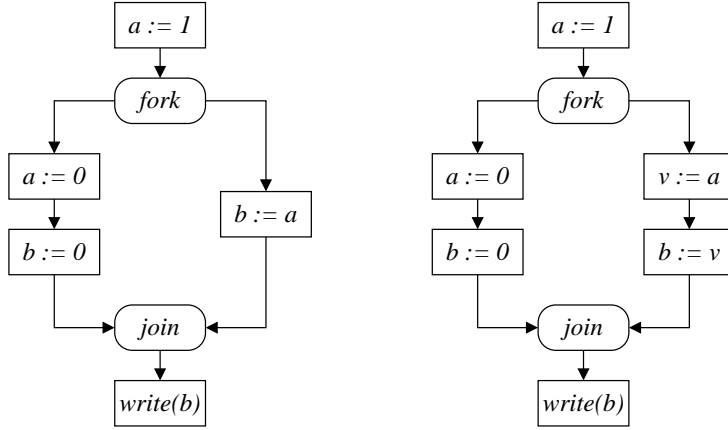


Fig. 6. Introduction of a virtual variable.

3.2 A Motivating Example

Consider the following program for which a control flow graph-like representation is shown in Figure 6, left:

$$a := 1; [(a := 0; b := 0) \parallel b := a]; \mathbf{write}(b).$$

Assume first that assignment statements execute atomically. It is not hard to show that under this assumption variable b is a (copy) constant of value 0 at the write instruction. In order to see this, note that in any execution $b := 0$ must be executed either after or before $b := a$ in the parallel thread. If it is executed after $b := a$ then b holds 0 at the write statement because the value 0 is assigned to b in the last executed assignment, $b := 0$. On the other hand, if $b := 0$ is executed before $b := a$ then also the reinitialization of a , $a := 0$, must have been executed before $b := a$ such that $b := a$ also loads the value 0 to b .

The situation is dramatically different, if assignment statements may execute non-atomically. In particular, if the assignment $b := a$ in the second thread is executed non-atomically, the first thread may execute the two statements $a := 0$ and $b := 0$ that kill a and b after a is loaded from the common memory but before the loaded value is stored to b . This results in a run of the program that propagates the value 1 from the initialization $a := 1$ to the final write-statement.

As explained in the previous section, we may model the two stage non-atomic execution of $b := a$ by splitting it into two assignments $v := a$ and $b := v$, where v is a new virtual variable that cannot be accessed by the first thread (cf. Figure 6, right). Because of this we can consider each of the virtual assignments $v := a$ and $b := v$ to be atomic. The resulting program has the run

$$\langle a := 1, v := a, b := 0, c := 0, b := v, \mathbf{write}(b) \rangle,$$

which—as the reader will have no difficulties to verify—propagates the value 1 from the initialization $a := 1$ to the write-statement. Thus, this run witnesses that b is not a copy constant at the write statement, in sharp contrast to the state of affairs under the assumption that assignments execute atomically.

3.3 The Domain of Non-Atomic Run Sets

A (*non-atomic*) run r is a sequence of assignments from the set \mathbf{Asg} defined above: $\mathbf{Runs} = \mathbf{Asg}^*$. We write $\mathbf{virtual}(r)$ for the set of virtual variables appearing in run r . As the specific choice of virtual variables is immaterial, we assume that all considered sets of runs are closed under bounded renaming of virtual variables. This enables a simple and adequate definition of the composition operators. In order to allow a technically clean treatment of this assumption, let $\equiv \subseteq \mathbf{Runs} \times \mathbf{Runs}$ be the equality of runs up to bounded renaming of virtual variables, i.e. $r \equiv r'$ hold if and only if r' can be obtained from r by bounded renaming of virtual variables.

Proposition 2 \equiv is an equivalence. \square

For a set of runs $R \subseteq \mathbf{Runs}$ we write R^\equiv for the closure of R w.r.t. \equiv :

$$R^\equiv = \{r \in \mathbf{Runs} \mid \exists r' \in R : r \equiv r'\}.$$

Obviously, this defines a closure operator.

Proposition 3

- (1) $R \subseteq R^\equiv$.
- (2) $(R^\equiv)^\equiv = R^\equiv$.
- (3) $R \subseteq S$ implies $R^\equiv \subseteq S^\equiv$. \square

The domain \mathbf{NR} is given by the sets of runs that are closed under \equiv :

$$\mathbf{NR} = \{R \subseteq \mathbf{Runs} \mid R = R^\equiv\}.$$

The members of \mathbf{NR} model sets of runs in a scenario where assignments execute non-atomically.

Lemma 4 (\mathbf{NR}, \subseteq) is a complete lattice with least element $\perp_{\mathbf{NR}} = \emptyset$ and greatest element $\top_{\mathbf{NR}} = \mathbf{Runs}$.

PROOF. (\mathbf{NR}, \subseteq) is a sub-lattice of the power set lattice $(2^{\mathbf{Runs}}, \subseteq)$. To show this, we have to check, that \mathbf{NR} is closed under arbitrary intersections and unions.

Here is the proof for intersection. Suppose $\mathcal{R} \subseteq \mathbf{NR}$ and $r, r' \in \mathbf{Runs}$ with $r \equiv r'$. We have to show that $r \in \bigcap \mathcal{R}$ if and only if $r' \in \bigcap \mathcal{R}$ which is simple:

$$\begin{aligned}
& r \in \bigcap \mathcal{R} \\
\text{iff} & \quad [\text{Definition of } \bigcap \mathcal{R}] \\
& \quad \forall R \in \mathcal{R} : r \in R \\
\text{iff} & \quad [\mathcal{R} \subseteq \mathbf{NR}, \text{ hence all } R \in \mathcal{R} \text{ are closed under } \equiv] \\
& \quad \forall R \in \mathcal{R} : r' \in R \\
\text{iff} & \quad [\text{Definition of } \bigcap \mathcal{R}] \\
& \quad r' \in \bigcap \mathcal{R}.
\end{aligned}$$

The proof for unions is just as simple and, therefore, omitted.

The least and greatest element of $(2^{\mathbf{Runs}}, \subseteq)$ are \emptyset and \mathbf{Runs} , respectively. It is obvious that both of them are closed under \equiv and hence are also the least and greatest elements, respectively, of (\mathbf{NR}, \subseteq) . \square

In the sections that follow we provide definitions for the operators and constants appearing in the constraint systems and show their well-definedness.

3.3.1 Base Statements

We can work with various atomicity assumptions as discussed above. The most natural and conservative one is that just single reads and writes of variables are atomic and that variables appearing more than once in an expression are re-read for every occurrence. This is captured by defining the semantics of an assignment statement, $\llbracket x := e(y_1, \dots, y_k) \rrbracket \in \mathbf{NR}$, where y_1, \dots, y_k refer to the variable occurrences in e , as the set of runs of the form

$$\langle v_{\pi(1)} := y_{\pi(1)}, \dots, v_{\pi(k)} := y_{\pi(k)}, x := e(v_1, \dots, v_k) \rangle,$$

where π is a permutation of $\{1, \dots, k\}$ and v_1, \dots, v_k are arbitrary distinct virtual variables. It is readily verified that $\llbracket x := e(y_1, \dots, y_k) \rrbracket$ is well-defined, i.e., that it is a member of \mathbf{NR} . We have to show that $\llbracket x := e(y_1, \dots, y_k) \rrbracket$ is closed under \equiv which is obvious as we admitted an arbitrary choice of virtual variables.

We may also work with more coarse-grained semantics of assignments. For our purposes the choice is arbitrary, as the dependence trace abstraction of an assignment will be precise with respect to any of these definitions.

Obviously, the only non-atomic run of statement **skip** is the empty run. Hence, $\llbracket \mathbf{skip} \rrbracket = \{\varepsilon\}$. Obviously, $\llbracket \mathbf{skip} \rrbracket \in \mathbf{NR}$.

The non-atomic runs induced by a base edge $e \in \mathbf{Base}$ are the non-atomic runs of the statement associated with e : $\llbracket e \rrbracket = \llbracket A(e) \rrbracket$, where $A(e)$ is the base statement associated with base edge e in the underlying flow graph.

3.3.2 Sequential Composition

The *sequential composition operator*, $\cdot; \cdot : \mathbf{NR} \times \mathbf{NR} \rightarrow \mathbf{NR}$, which is written as an infix operator, is defined by

$$R; S = \{r \cdot s \mid r \in R, s \in S, \mathbf{virtual}(r) \cap \mathbf{virtual}(s) = \emptyset\}^{\equiv}.$$

The condition about the local variables ensures that runs composed sequentially do not interact on local variables. The outer closure operator ensures that $;$ is well-defined

3.3.3 Interleaving Operator

The *interleaving operator*, $\otimes : \mathbf{NR} \times \mathbf{NR} \rightarrow \mathbf{NR}$, which we write in an infix form, is defined by

$$R \otimes S = \{r \mid \exists I_R, I_S : I_R \cup I_S = \{1, \dots, |r|\}, I_R \cap I_S = \emptyset, \\ r|_{I_R} \in R, r|_{I_S} \in S, \mathbf{virtual}(r|_{I_R}) \cap \mathbf{virtual}(r|_{I_S}) = \emptyset\}^{\equiv}.$$

The condition about the local variables in $r|_{I_R}$ and $r|_{I_S}$ ensures that parallel threads do not exchange values via local variables. The application of the closure operator $(\cdot)^{\equiv}$ guarantees well-definedness: $R \otimes S \in \mathbf{NR}$ for $R, S \in \mathbf{NR}$.

Suppose $r, s, t \in \mathbf{Runs}$ with $\mathbf{virtual}(r) \cap \mathbf{virtual}(s) = \emptyset$. We call t an *interleaving of r and s* if

$$\exists I_r, I_s : I_r \cup I_s = \{1, \dots, |r|\}, I_r \cap I_s = \emptyset, t|_{I_r} = r, t|_{I_s} = s$$

and denote the set of interleavings of r and s by $r \otimes s$.

3.3.4 Pre-Operator

The pre-operator, $pre : \mathbf{NR} \rightarrow \mathbf{NR}$ is defined as follows:

$$pre(R) = \{r \in \mathbf{Runs} \mid \exists r' \in \mathbf{Runs} : r \cdot r' \in R\}.$$

Lemma 5 *pre is well-defined.*

PROOF. We have to show that, for any $R \in \text{NR}$, $\text{pre}(R)$ is closed under \equiv . So suppose given $r, s \in \text{Runs}$ with $s \equiv r \in \text{pre}(R)$. Then there is $r' \in \text{Runs}$ with $r \cdot r' \in R$. By bounded renaming of local variables in r' we can construct a run s' such that $s \cdot s' \equiv r \cdot r'$. As R is closed under \equiv , $s \cdot s' \in R$ and hence $s \in \text{pre}(R)$. \square

3.3.5 Post-Operator

Analogously to the pre-operator, the post operator $\text{post} : \text{NR} \rightarrow \text{NR}$ is defined as follows:

$$\text{post}(R) = \{r \in \text{Runs} \mid \exists r' \in \text{Runs} : r' \cdot r \in R\}.$$

Lemma 6 *post is well-defined.* \square

3.4 Discussion

We have defined a complete lattice (NR, \subseteq) the members of which model sets of runs in a scenario in which assignment statements execute non-atomically. In order to enable an interleaving semantics to adequately capture the effect of non-atomic execution of assignments, we resorted to *virtual variables* that model storage locations that are private to threads. The members of NR are those sets of runs that are closed under bounded renaming of virtual variables. We have provided definitions for the operators and constants appearing in the constraint systems of Section 2. The (smallest) solution of these constraint systems over this new interpretation induces a semantics of parallel flow graphs that captures non-atomic execution of assignments. Thus, it provides another reference point for assessing flow analyses. This is put to advantage in the next section where we show that the dependence-traces interpretation developed there is a precise abstraction of the non-atomic interpretation of parallel flow graphs.

4 Dependence Traces

We can indirectly detect copy constants and eliminate faint code on the basis of the following information: given a program point u and a variable x of interest; when control is at another program point v , which variables y may

influence the value of x at u ? Clearly, this information can be derived from the set of bridging runs from u to v and we have a constraint system characterizing this set (cf. Section 2.5). We would like to compute the above information by means of a precise and effective abstract interpretation.

In this section, we develop an adequate abstract domain and adequate abstract operations for this. Our development will be guided by the requirements this domain must satisfy: (1) it must allow us to infer the above information easily; (2) it must allow us to define abstract operations that mirror precisely the corresponding operations on sets of (non-atomic) runs; and (3) it must allow us to compute fixpoints effectively.

Let us start with some definitions.

4.1 Transparency and Dependences

A run r is called *transparent* for a variables a if it does not contain an assignment with a as left hand side variable. Thus, a run is transparent for a if its execution is guaranteed not to change the value held by a .

Example 7 *The run $\langle a := 0, b := c \rangle$ is transparent for all variables except a and b , in particular for c . \square*

A *dependence* is a pair $d = (x, y)$ of program variables $x, y \in X$. We call x the *source variable* and y the *destination variable* of d . A run r is said to *exhibit* dependence (x, y) , if there are variables a_0, \dots, a_l , $l > 0$, expressions e_1, \dots, e_l , and (sub-) runs r_0, \dots, r_l such that

- (1) $r = r_0 \cdot \langle a_1 := e_1 \rangle \cdot r_1 \cdot \langle a_2 := e_2 \rangle \cdot r_2 \cdot \dots \cdot \langle a_l := e_l \rangle \cdot r_l$;
- (2) $a_0 = x$, $a_l = y$;
- (3) e_i contains a_{i-1} for $i = 1, \dots, l$; and
- (4) r_i is transparent for a_i for $i = 0, \dots, l$.

We also say “ (x, y) is a dependence of r ” in this case.

Example 8 *The run $\langle b := a, c := b, e := 0, f := e \rangle$ exhibits the dependences (a, b) , (a, c) , and (b, c) but not the dependence (e, f) because e is killed by the assignment $e := 0$ before it is read. \square*

4.2 Dependence Traces

Unfortunately, we cannot use dependences themselves as abstract domain because, in general, the dependences of the interleavings of two runs (or runs

sets) cannot directly be inferred from the dependences of these runs (or run sets). As an example, consider the two runs $r_1 = \langle b := a, b := 0, d := c \rangle$ and $r_2 = \langle d := c \rangle$. Both exhibit just the dependence (c, d) . But r_1 can be interleaved with $r_3 = \langle c := b \rangle$ to $\langle b := a, c := b, b := 0, d := c \rangle$, a run that exhibits (a, d) while r_2 cannot. Thus, an abstraction of run sets that faithfully mirrors dependences must collect more information than just dependences. We propose to employ *dependence traces* that are defined in the following.

The basic idea is to collect not only dependences but sequences of dependences that can successively be exhibited by a run. For example, we record the sequence $\varphi = \langle (a, b), (c, d) \rangle$ for the run r_1 but not for r_2 . Intuitively, φ plays a dual role: on the one hand, it captures the potential of r_1 to exhibit dependence (a, d) if a run of a parallel thread fills the gap between b and c (like run r_3) and, on the other hand, its potential to successively fill the gaps (a, b) and (c, d) in a parallel run (like in $\langle a := x, c := b, y := d \rangle$).

A *dependence sequence* is a sequence $\varphi = \langle (x_1, y_1), \dots, (x_k, y_k) \rangle$, $k \geq 0$, of dependences. Note that we allow the empty dependence sequence ε . We write $\overleftarrow{\varphi}$ for x_1 and $\overrightarrow{\varphi}$ for y_k , if $\varphi \neq \varepsilon$; if $\varphi = \varepsilon$, $\overleftarrow{\varphi}$ and $\overrightarrow{\varphi}$ are undefined. We denote the set of dependence sequences by **DS**.

Example 9 $\varphi = \langle (a, b), (c, d) \rangle$ is a dependence sequence with $\overleftarrow{\varphi} = a$ and $\overrightarrow{\varphi} = d$. \square

Further information must be collected. To see why, compare the run $r_4 = \langle a := 0, b := a, b := 0, d := c \rangle$ with r_1 . Unlike r_1 , r_4 does *not* have the potential to exhibit dependence (a, d) if a parallel run fills the gap between b and c , but like r_1 it can successively fill the gaps (a, b) and (c, d) in a parallel run. The crucial difference is that in r_4 the part of the run before a is read is not transparent for the source variable of the first dependence, viz. a . A similar difference can arise for the target variable of the final dependence. Therefore, we refine dependence sequences to *dependence traces*.

A *dependence trace* is a triple $\tau = (\iota, \varphi, \kappa)$ consisting of Boolean values $\iota, \kappa \in \mathbb{B} = \{0, 1\}$ coding initial and final transparency and a dependence sequence φ . We assume that $\iota = 0$ and $\kappa = 0$ if $\varphi = \varepsilon$. The set of dependence traces is denoted by **DT**:

$$\mathbf{DT} = \{(\iota, \varphi, \kappa) \in \mathbb{B} \times \mathbf{DS} \times \mathbb{B} \mid \varphi = \varepsilon \Rightarrow (\iota = 0 \wedge \kappa = 0)\}.$$

A run r is said to *exhibit* dependence trace $\tau = (\iota, \langle (x_1, y_1), \dots, (x_k, y_k) \rangle, \kappa)$, $r \vdash \tau$ for short, if there are sub-runs $t_0, \dots, t_k, r_1, \dots, r_k$, such that

- (1) $r = t_0 \cdot r_1 \cdot t_1 \cdot r_2 \cdots r_k \cdot t_k$;
- (2) r_i exhibits dependence (x_i, y_i) for $i = 1, \dots, k$;
- (3) $\iota = 1$ implies that t_0 is transparent for x_1 ; and

(4) $\kappa = 1$ implies that t_k is transparent for y_k .

In this case, we call $t_0 \cdot r_1 \cdot t_1 \cdot r_2 \cdots r_k \cdot t_k$ a *decomposition* of r that witnesses $r \vdash \tau$. Note that $r \vdash (0, \varepsilon, 0)$ holds for all runs r as witnessed by the trivial decomposition $t_0 = r$. The trivial dependence trace $(0, \varepsilon, 0)$ allows us to distinguish the dependence trace abstraction of the empty run set from the abstraction of non-empty run sets.

Instead of saying “ r exhibits τ ” we often use the phrase “ τ is a dependence trace of r ”.

Example 10 Run r_1 exhibits the dependence trace $(1, \langle (a, b), (c, d) \rangle, 1)$ in contrast to r_4 . However, both runs share the dependence trace $(0, \langle (a, b), (c, d) \rangle, 0)$. \square

Example 11 Consider the run $r_5 = \langle a := 0, b := a, c := b, c := 0, f := e, e := 0 \rangle$. One of the dependence traces of r_5 is $\tau = (0, \langle (a, c), (e, f) \rangle, 1)$ as witnessed by the decomposition $r = t_0 \cdot r_1 \cdot t_1 \cdot r_2 \cdot t_2$ where

$$\langle \underbrace{a := 0}_{t_0}, \underbrace{b := a, c := b}_{r_1}, \underbrace{c := 0}_{t_1}, \underbrace{f := e}_{r_2}, \underbrace{e := 0}_{t_2} \rangle.$$

Another decomposition witnessing τ is

$$\langle \underbrace{a := 0}_{t_0}, \underbrace{b := a, c := b}_{r_1}, \underbrace{c := 0}_{t_1 = \varepsilon}, \underbrace{f := e, e := 0}_{r_2}, \underbrace{}_{t_2 = \varepsilon} \rangle.$$

Run r_5 has also many other dependence traces, e.g., $(1, \langle (b, c), (e, f) \rangle, 1)$ and $(1, \langle (e, f) \rangle, 1)$. \square

Ultimately, we are interested in dependence traces without gaps that code complete transfers from one variable to another one, where a gap can either be a lack of initial or final transparency or a hole from y_i to x_{i+1} . Thus, the dependence traces of ultimate interest are those of the form $(1, \langle (x, y) \rangle, 1)$. They correspond to dependences.

Proposition 12 $r \vdash (1, \langle (x, y) \rangle, 1)$ if and only if r exhibits dependence (x, y) . \square

The abstraction of run sets must allow us to propagate the transparency bit of dependence traces through sequential contexts. For this purpose we collect in addition to a set of dependence traces a set of variables for which a transparent run exists. According to these ideas, we may abstract a set of (non-atomic) runs R to a pair (T_R, D_R) consisting of the set of variables

$$T_R := \{x \mid \exists r \in R : r \text{ is transparent for } x\}$$

and the set of dependence traces

$$D_R := \{\tau \mid \exists r \in R : r \vdash \tau\}.$$

Indeed, on this abstraction of run sets, we can define abstract operators that precisely mirror the operators on sets of non-atomic runs. Still we are not yet done. The problem is that we do not know how to represent D_R finitely.

Fortunately, it is not necessary to collect *all* dependence traces in the abstraction, in order to describe the potential for forming dependences with a parallel context. It suffices to retain only certain “short” dependence traces in the abstraction that subsume the potential of all the other ones. Before we turn to the technical development, let us illustrate this kind of subsumption by a small example.

Consider the two dependence traces $\tau_1 = (1, \langle (a, b), (c, d), (e, f) \rangle, 1)$ and $\tau_2 = (1, \langle (a, d), (e, f) \rangle, 1)$. Both share the gap (d, e) but τ_1 has the additional gap (b, c) . If a run of a parallel context can successively fill the two gaps in τ_1 —i.e. if it exhibits the dependence trace $\tau_3 = (0, \langle (b, c), (d, e) \rangle, 0)$ —it can also fill the single gap in τ_2 —i.e. it also exhibits $\tau_4 = (0, \langle (d, e) \rangle, 0)$. Two interesting relationships between dependence traces popped up in this discussion. On the one hand, τ_1 is “subsumed” by τ_2 . On the other hand τ_4 is “implied” by τ_3 as it has less dependences: any run having τ_3 as a dependence traces also has τ_4 as a dependence trace. We now define two orders on the set of dependence traces that capture these two relationships, the “implication order” and the “subsumption order”.

4.3 Implication Order

Let $\leq \subseteq \text{DT} \times \text{DT}$ be the smallest reflexive and transitive relation on the set of dependence traces that satisfies

- (1) $(\iota, \varphi \cdot \langle (x, y) \rangle \cdot \psi, \kappa) \leq (\iota, \varphi \cdot \psi, \kappa)$, if $\varphi \neq \varepsilon \vee \iota = 0$ and $\psi \neq \varepsilon \vee \kappa = 0$;
- (2) $(1, \varphi, \kappa) \leq (0, \varphi, \kappa)$; and
- (3) $(\iota, \varphi, 1) \leq (\iota, \varphi, 0)$.

Proposition 13 \leq is a partial order on DT called the implication order. \square

The implication order \leq allows us to weaken the information in a dependence trace in two ways. First of all, we can omit dependences (1); here we must be careful not to omit the first or last dependence if the corresponding transparency bit is set, as otherwise the transparency bit might become invalid. Secondly, we can weaken the information about transparency of the initial or final part of the run, by changing the transparency bits from 1 to 0 (2 & 3).

The most appealing fact about \leq is that it preserves compatibility, which justifies the name “implication order”.

$$\begin{array}{rcl}
\tau: & a \underline{\quad} b & c \underline{\quad} d & e \underline{\quad} f & g \underline{\quad} h \\
\tau \leq \tau': & a \underline{\quad} b & & e \underline{\quad} f & g \underline{\quad} h \\
\tau \sqsubseteq \tau'': & a \underline{\quad} b & c \underline{\quad\quad\quad} & f & g \underline{\quad} h
\end{array}$$

Fig. 7. Illustration of implication and subsumption order.

Proposition 14 (\leq preserves compatibility) *Suppose $r \vdash \tau$ and $\tau \leq \tau'$. Then $r \vdash \tau'$. \square*

Example 15 *Consider the dependence trace $\tau = (1, \langle(a, b), (c, d)\rangle, 0)$ of the run $r = \langle b := a, c := 0, d := c, d := 0 \rangle$. Here is a list of the dependence traces implied by τ :*

$$\begin{aligned}
\tau_1 &= (0, \langle(a, b), (c, d)\rangle, 0) \\
\tau_2 &= (1, \langle(a, b)\rangle, 0) \\
\tau_3 &= (0, \langle(a, b)\rangle, 0) \\
\tau_4 &= (0, \langle(c, d)\rangle, 0) \\
\tau_5 &= (0, \varepsilon, 0)
\end{aligned}$$

i.e., we have $\tau \leq \tau_i$ for $i = 1, \dots, 5$. All of them are dependence traces of r . But we do not have $\tau \leq \tau_6$ for $\tau_6 = (1, \langle(c, d)\rangle, 0)$. And indeed, τ_6 is not a dependence trace of r because variable c is killed before it is read in r . \square

4.4 Subsumption Order

We now define the *subsumption order* $\sqsubseteq \subseteq \text{DT} \times \text{DT}$. Intuitively, $\tau \sqsubseteq \tau'$ captures that τ' has fewer gaps than τ and thus subsumes the potential of τ for forming dependences with a cooperating parallel context. We define \sqsubseteq as the smallest transitive and reflexive relation that satisfies

$$(\iota, \varphi \cdot \langle(x, y)\rangle \cdot \varphi' \cdot \langle(x', y')\rangle \cdot \varphi'', \kappa) \sqsubseteq (\iota, \varphi \cdot \langle(x, y')\rangle \cdot \varphi'', \kappa).$$

Fig. 7 illustrates the difference between the implication and the subsumption order. For simplicity, we only show the dependence sequences and omit the transparency bits. In the top row we show a dependence trace τ , in the middle row a dependence trace τ' that is implied by τ , and in the bottom row a dependence trace τ'' that subsumes τ . The implication order allows us to omit dependences (and weaken transparency bits). In contrast the subsumption order allows us to remove gaps.

It is obvious from the defining rule that a dependence trace τ' that properly subsumes another dependence trace τ embodies a strictly shorter dependence sequence. Therefore, \sqsubseteq satisfies the ascending chain condition.

Proposition 16 \sqsubseteq is a partial order on DT that satisfies the ascending chain condition: every strictly increasing sequence $\tau_1 \sqsubset \tau_2 \sqsubset \dots$ is finite. \square

Note that dependence traces of the form $(1, \langle(x, y)\rangle, 1)$, which correspond to dependences by Proposition 12, are maximal w.r.t. \sqsubseteq . This simple observation is important, as it implies that we cover all dependences even when we only consider \sqsubseteq -maximal dependence traces.

4.5 A Lattice of Antichains

An *antichain* with respect to \sqsubseteq (or \sqsubseteq -antichain for short) is a set $D \subseteq \text{DT}$ of dependence traces satisfying

$$\neg \exists \tau, \tau' \in D : \tau \sqsubset \tau'.$$

We denote the set of \sqsubseteq -antichains by AC. We can lift the subsumption order to AC as follows:

$$D \sqsubseteq D' \quad :\Leftrightarrow \quad \forall \tau \in D \exists \tau' \in D' : \tau \sqsubseteq \tau'.$$

Thus, D' subsumes D , if every dependence trace in D is subsumed by some dependence trace in D' . We call \sqsubseteq the *antichain order*. This is justified by the following lemma.

Lemma 17 \sqsubseteq is a partial order on AC.

PROOF. It is straightforward to show that \sqsubseteq is reflexive and transitive. Let us show that \sqsubseteq is also antisymmetric and hence a partial order.

Suppose $D \sqsubseteq D' \sqsubseteq D$. We show that $D \subseteq D'$, the reverse inclusion follows analogously. Suppose $\tau \in D$. Then there is $\tau' \in D'$ with $\tau \sqsubseteq \tau'$ as $D \sqsubseteq D'$. Because of $D' \sqsubseteq D$, there is $\tau'' \in D$ with $\tau' \sqsubseteq \tau''$. Thus, we have

$$D \ni \tau \sqsubseteq \tau' \sqsubseteq \tau'' \in D.$$

As D is an antichain, this implies that $\tau = \tau''$. Consequently, all these three dependence traces must be equal: $\tau = \tau' = \tau''$. But then $\tau = \tau' \in D'$. \square

A simple way to form an \sqsubseteq -antichain out of an arbitrary subset $D \subseteq \text{DT}$ is to consider the set of \sqsubseteq -maximal elements in D . We denote this set by D^\uparrow :

$$D^\uparrow = \{\tau \in D \mid \neg \exists \tau' \in D : \tau \sqsubset \tau'\}.$$

The dependence traces in D^\uparrow subsume all dependence traces in D . In this sense, no interesting information is lost when going from D to D^\uparrow .

Lemma 18 (\uparrow **subsumes**) *For any $\tau \in D$ there is a $\tau' \in D^\uparrow$ such that $\tau \sqsubseteq \tau'$.*

PROOF. The lemma follows easily with the ascending chain condition. \square

The operator \uparrow is a co-closure operator that yields \sqsubseteq -antichains:

Lemma 19 (\uparrow **is a co-closure operator**)

- (1) $D^\uparrow \subseteq D$.
- (2) $(D^\uparrow)^\uparrow = D^\uparrow$.
- (3) D^\uparrow is an \sqsubseteq -antichain.
- (4) $(\cdot)^\uparrow$ is monotonic: $D \subseteq E$ implies $D^\uparrow \subseteq E^\uparrow$.

PROOF. The proof of these properties is straightforward. \square

The \sqsubseteq -antichains together with the lifted subsumption order form a complete lattice.

Lemma 20 (AC, \sqsubseteq) *is a complete lattice. The least upper bound (lub) of a subset $\mathcal{D} \subseteq \text{AC}$ is $\bigsqcup \mathcal{D} := (\bigcup \mathcal{D})^\uparrow$ and the least element of (AC, \sqsubseteq) is $\perp_{\text{AC}} := \emptyset$.*

PROOF. In order to show that (AC, \sqsubseteq) is a complete lattice, it suffices to demonstrate that any subset $\mathcal{D} \subseteq \text{AC}$ has a least upper bound. We show that, as claimed in the lemma, $E := (\bigcup \mathcal{D})^\uparrow$ is indeed the least upper bound of \mathcal{D} .

Firstly, E is an upper bound of \mathcal{D} : we have to show that $D \sqsubseteq E$ for any $D \in \mathcal{D}$, which is seen as follows:

$$\begin{aligned}
& \tau \in D \\
\Rightarrow & [D \in \mathcal{D}, \text{definition of } \bigcup \mathcal{D}] \\
& \tau \in \bigcup \mathcal{D} \\
\Rightarrow & [\text{Lemma 18, definition } E] \\
& \exists \tau' \in E : \tau \sqsubseteq \tau'.
\end{aligned}$$

Secondly, E is smaller than any other bound \mathcal{D} . Suppose F is an arbitrary upper bound of \mathcal{D} . Then $E \sqsubseteq F$ follows from the following chain of implications:

$$\begin{aligned}
& \tau \in E \\
\Rightarrow & \quad [\text{Definition } E, \text{ Lemma 19(1.)}] \\
& \tau \in \bigcup \mathcal{D} \\
\Rightarrow & \quad [\text{Definition of } \bigcup \mathcal{D}] \\
& \exists D \in \mathcal{D} : \tau \in D \\
\Rightarrow & \quad [D \sqsubseteq F \text{ as } F \text{ is an upper bound of } \mathcal{D}, \text{ definition } \sqsubseteq] \\
& \exists \tau' \in F : \tau \sqsubseteq \tau'.
\end{aligned}$$

The least element of $(\mathbf{AC}, \sqsubseteq)$ is $\perp_{\mathbf{AC}} = \bigsqcup \emptyset = (\emptyset)^\uparrow = \emptyset$. \square

Let us consider another operator on sets of dependence traces, the downwards closure operator $(\cdot)^\downarrow$. It is defined for sets $D \in \mathbf{DT}$ by

$$D^\downarrow = \{\tau \in \mathbf{DT} \mid \exists \tau' \in D : \tau \sqsubseteq \tau'\}.$$

We can apply $(\cdot)^\downarrow$ in particular to antichains. Thus, we may consider $(\cdot)^\downarrow$ as an operator $(\cdot)^\downarrow : \mathbf{AC} \rightarrow 2^{\mathbf{DT}}$. It is not hard to see that $(\cdot)^\downarrow$ is monotonic.

Proposition 21 *Suppose $A, B \in \mathbf{AC}$. Then $A \sqsubseteq B$ implies $A^\downarrow \subseteq B^\downarrow$. \square*

$(\cdot)^\uparrow$ and $(\cdot)^\downarrow$ are approximate inverses of each other.

Lemma 22 *For any $D \in \mathbf{DT}$, we have $D^{\uparrow\downarrow} \supseteq D$ and $D^{\downarrow\uparrow} \subseteq D$. For any $A \in \mathbf{AC}$, we even have $A^{\downarrow\uparrow} = A$. As a consequence, $((\cdot)^\uparrow, (\cdot)^\downarrow)$ is a Galois surjection from $2^{\mathbf{DT}}$ to \mathbf{AC} :*

$$2^{\mathbf{DT}} \begin{array}{c} \xrightarrow{(\cdot)^\uparrow} \\ \rightleftarrows \mathbf{AC} \\ \xleftarrow{(\cdot)^\downarrow} \end{array}$$

PROOF.

$D^{\uparrow\downarrow} \supseteq D$: By Lemma 18, there is, for any $\tau \in D$, a dependence trace $\tau' \in D^\uparrow$ such that $\tau \sqsubseteq \tau'$. This implies that $\tau \in D^{\uparrow\downarrow}$.

$D^{\downarrow\uparrow} \subseteq D$: If $\tau \in D^{\downarrow\uparrow}$, then τ is a maximal element in D^\downarrow . The maximal elements in D^\downarrow , however, must already be in D , as they cannot be added to D by lying strictly below another element of D .

$A^{\downarrow\uparrow} = A$: It remains to show that $A^{\downarrow\uparrow} \supseteq A$. Any $\tau \in A$ is maximal in A^\downarrow . Therefore, any such τ is also in $A^{\downarrow\uparrow}$. \square

The fact that $((\cdot)^\uparrow, (\cdot)^\downarrow)$ is a Galois surjection from 2^{DT} into AC shows us that \sqsubseteq -antichains form a reasonable abstraction of sets of dependence traces. It also has other interesting consequences. First of all, it implies that $(\cdot)^\uparrow$ is universally disjunctive, which is important for ensuring that the abstraction mapping and the abstract operators defined later are universally disjunctive as well.

Proposition 23 $(\cdot)^\uparrow : 2^{\text{DT}} \rightarrow \text{AC}$ is universally disjunctive (*‘distributive’*). \square

Secondly, it shows us that we can present (AC, \sqsubseteq) isomorphically by downwards closed sets of dependence traces. From the theory of Galois connections, we know that the images of the upper and lower adjoint are isomorphic. This implies that (AC, \sqsubseteq) , the image of $(\cdot)^\uparrow$, is isomorphic to the image of $(\cdot)^\downarrow$, which is the set of downwards closed sets of dependence traces ordered by set inclusion. Note that this isomorphism depends on the fact that the underlying subsumption order on dependence traces satisfies the ascending chain condition. Otherwise, Lemma 18 would fail and we would not have the property $D^{\uparrow\downarrow} \supseteq D$ that is crucial for the isomorphism between antichains and downwards closed sets.

For our purpose it is more convenient to work with antichains, because this leads to a more natural definition of the interleaving operator. If we work with downwards closed sets we may add dependence traces by means of downwards closure that are not exhibited by any run in the abstracted run set. These additional dependence traces do not represent actual potential of the run set and in order to avoid imprecision, we must ensure that they are not considered for inferring dependence traces of interleavings.

4.6 Short Dependence Traces

A dependence sequence $\varphi = \langle (x_1, y_1), \dots, (x_k, y_k) \rangle$ is called *short* if

- (1) all destination variables of dependences not counting the last one are distinct: for all $1 \leq i < j < k$, $y_i \neq y_j$; and
- (2) all source variables of dependences not counting the first one are distinct: for all $1 < i < j \leq k$, $x_i \neq x_j$.

A dependence trace $\tau = (\iota, \varphi, \kappa)$ is called *short* if the embodied dependence sequence φ is short. We write DTS for the set of short dependence traces:

$$\text{DTS} = \{ \tau \in \text{DT} \mid \tau \text{ is short} \}.$$

Example 24 Consider the run $r = \langle c := a, c := b, e := d \rangle$. One of its dependence traces is $\tau = (1, \langle (a, c), (b, c), (d, e) \rangle, 1)$, which is not short due to the

repetition of variable c as a target variable. But run r has also the dependence trace $\tau' = (1, \langle (a, c), (d, e) \rangle, 1)$ which is short and subsumes τ . This is not a coincidence as we will see in a moment (Lemma 26). \square

We are interested in short dependence traces for two reasons. Firstly, there are only finitely many short dependence traces. This makes the abstract domain introduced in the next section finite as well and ensures that fixpoints for monotonic functions on this domain can be calculated effectively. The following lemma provides a formula for the cardinality of DTS and an asymptotic bound.

Lemma 25 *Let $n = |X|$. Then $|\text{DTS}| = 1 + 4n^2 n!^2 \sum_{i=0}^n \frac{1}{i!^2} = \mathcal{O}(n^{2n+2})$.*

PROOF. By the pigeonhole principle, a dependence sequence cannot contain more than $n + 1$ dependences without violating the condition of shortness.

Let $i \in \{0, \dots, n\}$. For forming a dependence sequence $\langle d_0, \dots, d_i \rangle$ of length $i + 1$ in a dependence trace, we can choose arbitrary program variables as source variable of d_0 and as destination variable of d_i ; there are n^2 ways of doing this. We can choose the remaining source variables of d_1, \dots, d_i as an arbitrary i -permutation of the variables in X . (Recall that an i -permutation of X is an ordered sequence of i elements of X , with no element appearing more than once in the sequence). The same holds for the remaining destination variables of d_0, \dots, d_{i-1} . As there are $\frac{n!}{(n-i)!}$ i -permutations [28], there are thus $n^2 \left(\frac{n!}{(n-i)!}\right)^2$ short dependence sequences of length $i + 1$. There are four possible choices for the transparency bits in a dependence trace with a given non-empty dependence sequence. In addition we have a single dependence trace with an empty dependence sequence, viz. $(0, \varepsilon, 0)$. Summing up, the number of short dependence traces is thus

$$1 + 4 \sum_{i=0}^n \left(n^2 \left(\frac{n!}{(n-i)!} \right)^2 \right) = 1 + 4n^2 n!^2 \sum_{i=0}^n \frac{1}{(n-i)!^2} = 1 + 4n^2 n!^2 \sum_{i=0}^n \frac{1}{i!^2}.$$

Using the well-known fact that $n! \leq n^n$ and bounding the sum by

$$\sum_{i=0}^n \frac{1}{i!^2} \leq \sum_{i=0}^n \frac{1}{i!} \leq \sum_{i=0}^{\infty} \frac{1}{i!} = e$$

the asymptotic bound $\mathcal{O}(n^{2n+2})$ follows. \square

The asymptotic bound $\mathcal{O}(n^{2n+2})$ for $|\text{DTS}|$ is rather rough as it involves the rather bad estimate n^n for $n!$ but suffices for our purposes. Using for instance Stirling's approximation [28] for the factorial function, we could obtain tighter bounds.

The second reason why we are interested in short dependence traces is that they suffice to capture the potential of runs to aid in forming dependences ‘up to subsumption’ as the following lemma shows.

Lemma 26 (Short dependence traces subsume) *Suppose $r \vdash \tau$. Then there is a short dependence trace τ' with $r \vdash \tau'$ and $\tau \sqsubseteq \tau'$.*

PROOF. Suppose $r \vdash \tau = (\iota, \langle (x_1, y_1), \dots, (x_k, y_k) \rangle, \kappa)$. We describe a shortening procedure that can be iterated until a short dependence trace is obtained.

Suppose τ is not already short. Let us assume that Condition 1. is violated; if Condition 2. is violated we can proceed analogously. Then there are indices i, j , $1 \leq i < j < k$, with $y_i = y_j$. Consider the dependence trace τ' obtained from τ by removing the middle part $\langle (x_{i+1}, y_{i+1}), \dots, (x_j, y_j) \rangle$ of the dependence sequence:

$$\tau' := (\iota, \langle (x_1, y_1), \dots, (x_i, y_i), (x_{j+1}, y_{j+1}), \dots, (x_k, y_k) \rangle, \kappa).$$

It is not hard to see that both $\tau \sqsubseteq \tau'$ and $\tau \leq \tau'$. By Proposition 14 the latter implies $r \vdash \tau'$. \square

We still have to see that we can obtain the short dependence traces of a composed set of runs from the short dependence traces of the argument run sets. This is particularly challenging for run sets obtained by interleaving and will be the topic of Sections 4.8–4.12.

Shortening a dependence trace w.r.t. either \leq or \sqsubseteq results again in a short dependence trace.

Lemma 27 (\leq and \sqsubseteq preserve shortness) *If τ is short and $\tau \leq \tau'$ or $\tau \sqsubseteq \tau'$, then τ' is short.*

PROOF. All pairs of source or target variables in τ' are also pairs of target variables in τ if $\tau \leq \tau'$ or $\tau \sqsubseteq \tau'$. \square

We denote the set of antichains of short dependence traces by ACS:

$$\text{ACS} = \{D \in \text{AC} \mid D \subseteq \text{DTS}\}.$$

Lemma 26 implies that \sqsubseteq -maximal dependence traces of a run (or run set) are always short. Therefore, if we restrict attention to short dependence traces of a run or run set, we still capture all maximal dependence traces. By working

with ACS instead of AC, we code this knowledge into the domain. In particular, we do not lose dependences because the dependence traces of the form $(1, \langle (a, b) \rangle, 1)$ that correspond to dependences are trivially short.

Lemma 28 $(\text{ACS}, \sqsubseteq)$ is a complete sub-lattice of (AC, \sqsubseteq) . Its height is $|\text{DTS}| + 1 = \mathcal{O}(n^{2n+2})$ where $n = |X|$.

PROOF. Suppose $\mathcal{D} \subseteq \text{ACS}$. In order to prove that $(\text{ACS}, \sqsubseteq)$ is a complete sub-lattice of (AC, \sqsubseteq) we have to show that $\sqcup \mathcal{D} \in \text{ACS}$, i.e. that $\sqcup \mathcal{D} \subseteq \text{DTS}$:

$$\begin{array}{ccccc} \sqcup \mathcal{D} & = & (\cup \mathcal{D})^\uparrow & \subseteq & \cup \mathcal{D} & \subseteq & \text{DTS}. \\ & & \uparrow & & \uparrow & & \uparrow \\ & & [\text{Lem. 20}] & & [\text{Lem. 19}] & & [\mathcal{D} \subseteq \text{ACS}] \end{array}$$

We can restrict the downwards closure operator to short dependence traces, i.e. redefine it by $D^\downarrow = \{\tau \in \text{DTS} \mid \exists \tau' \in D : \tau \sqsubseteq \tau'\}$ for $D \subseteq \text{DTS}$. It follows as in Lemma 22 that $((\cdot)^\uparrow, (\cdot)^\downarrow)$ is a Galois surjection from 2^{DTS} into ACS:

$$2^{\text{DTS}} \begin{array}{c} (\cdot)^\uparrow \\ \xrightarrow{\quad} \\ (\cdot)^\downarrow \end{array} \text{ACS}$$

As a consequence $(\text{ACS}, \sqsubseteq)$ is isomorphic to the lattice of downwards closed subsets of DTS, ordered by set inclusion. The latter is a sub-lattice of $(2^{\text{DTS}}, \subseteq)$. Hence its height (and thus the height of $(\text{ACS}, \sqsubseteq)$) cannot be larger than the height of $(2^{\text{DTS}}, \subseteq)$ which is $|\text{DTS}| + 1$.

On the other hand, we can construct an ascending chain of size $|\text{DTS}| + 1$. Let $(x_1, \dots, x_{|\text{DTS}|})$ be a topological sort of $(\text{DTS}, \sqsubseteq)$, i.e., a list containing all elements of DTS such that $x_i \sqsubseteq x_j$ implies $i \leq j$ for all $i, j \in \{1, \dots, |\text{DTS}|\}$. Then we can define a chain of length $\text{DTS} + 1$ by choosing $A_0 = \emptyset$ and $A_i = (A_{i-1} \cup \{x_i\})^\uparrow$ for $i = 1, \dots, |\text{DTS}|$. $A_{i-1} \sqsubseteq A_i$ is obvious, and $A_{i-1} \neq A_i$ holds because $A_{i-1} \subseteq \{x_1, \dots, x_{i-1}\}$, which is seen by a straightforward induction. Thus, x_i is maximal in $A_{i-1} \cup \{x_i\}$ due to the topological sort property.

The asymptotic bound $|\text{DTS}| + 1 = \mathcal{O}(n^{2n+2})$ follows from Lemma 25. \square

4.7 The Abstract Domain

Let us now define the abstract domain. The values of the abstract domain AD are pairs (T, D) consisting of a set $T \subseteq X$ of variables and an \sqsubseteq -antichain D

of short dependence traces:

$$\text{AD} = 2^X \times \text{ACS}.$$

T represents the variables for which a transparent run exists. This information is necessary in order to allow a proper propagation of initial and final transparency bits in sequential contexts. The order on the abstract domain, which we also denote by the symbol \sqsubseteq , is defined as the lift of the inclusion order on the T component and the antichain order \sqsubseteq on the D component: $(T, D) \sqsubseteq (T', D')$ iff

- (1) $T \subseteq T'$ and
- (2) $D \sqsubseteq D'$.

(AD, \sqsubseteq) is the product lattice of the complete lattices $(2^X, \subseteq)$ and $(\text{ACS}, \sqsubseteq)$ and hence also a complete lattice. Both of these lattices have \emptyset as their least element. Hence, (\emptyset, \emptyset) is the least element of \sqsubseteq .

Lemma 29 *(AD, \sqsubseteq) is a complete lattice with least element (\emptyset, \emptyset) . Its height is $\mathcal{O}(n^{2n+2})$ where $n = |X|$.*

PROOF. It only remains to prove the asymptotic bound for the height. The height of AD is the sum of the height of $(2^X, \subseteq)$, which is $n + 1$, and the height of $(\text{ACS}, \sqsubseteq)$, which is $\mathcal{O}(n^{2n+2})$ by Lemma 28. This implies the stated bound. \square

Let us now define an abstraction mapping $\alpha : \text{NR} \rightarrow \text{AD}$ that captures the intuition how non-atomic run sets are abstracted to values from AD :

$$\begin{aligned} \alpha(R) &= (T_R, D_R), \text{ where} \\ T_R &= \{x \in X \mid \exists r \in R : r \text{ is transparent for } x\} \text{ and} \\ D_R &= \{\tau \in \text{DT} \mid \exists r \in R : r \vdash \tau\}^\uparrow. \end{aligned}$$

Before we proceed, let us show that this is a proper definition.

Lemma 30 *α is well-defined.*

PROOF. We have to show two things for an arbitrary $R \in \text{NR}$:

- (1) D_R consists of short dependence traces.
- (2) D_R is an \sqsubseteq -antichain.

To 1.: Assume there is $\tau \in D_R$ that is not short. Then there is $r \in R$ with $r \vdash \tau$. By Lemma 26, there is a *short* dependence trace τ' with $r \vdash \tau$ and $\tau \sqsubseteq \tau'$. In particular $\tau' \in \{\tau \in \text{DT} \mid \exists r \in R : r \vdash \tau\}$ and, as τ' is short and τ is not, we even have $\tau \sqsubset \tau'$. But this shows that τ is not maximal in $\{\tau \in \text{DT} \mid \exists r \in R : r \vdash \tau\}$ and hence is not a member of D_R , a contradiction.

To 2.: This is ensured by Lemma 19(4). \square

The abstraction $\alpha(R)$ of a run set R is induced by the following abstraction $\beta(r)$ of the single runs $r \in R$:

$$\begin{aligned} \beta(r) &= (T_r, D_r), \text{ where} \\ T_r &= \{x \in X \mid r \text{ is transparent for } x\} \text{ and} \\ D_r &= \{\tau \in \text{DT} \mid r \vdash \tau\}^\uparrow. \end{aligned}$$

Lemma 31 *Suppose $R \in \text{NR}$. Then $\alpha(R) = \sqcup\{\beta(r) \mid r \in R\}$.*

PROOF. We have $\sqcup\{\beta(r) \mid r \in R\} = (\bigcup_{r \in R} T_r, \sqcup_{r \in R} D_r)$. It is obvious that $T_R = \bigcup_{r \in R} T_r$. On the other hand, we have $\sqcup_{r \in R} D_r = (\bigcup_{r \in R} \{\tau \mid r \vdash \tau\}^\uparrow)^\uparrow$, by Lemma 20. It is not hard to show that this equals D_R by considering the \sqsubseteq - and the \sqsupseteq -direction separately. \square

The fact that α is induced by an abstraction on single runs has nice consequences.

Proposition 32 *α is monotonic: $R \subseteq R'$ implies $\alpha(R) \sqsubseteq \alpha(R')$.* \square

Proposition 33 *α is universally disjunctive.* \square

The latter property is crucial for precision of the abstract interpretation of constraint systems, cf. Section 5, and shows us that α provides a proper abstraction of run sets by being the lower adjoint of a Galois connection. For completeness let us introduce the corresponding upper adjoint. It is $\gamma : \text{AD} \rightarrow \text{NR}$, defined by

$$\gamma(T, D) = \{r \mid T_r \subseteq T, D_r \sqsubseteq D\}.$$

Proposition 34 *(α, γ) is a Galois connection between NR and AD : $\text{NR} \stackrel{\alpha}{\rightleftarrows} \text{AD}$.* \square

We leave the proof that γ is well-defined and forms a Galois connection with α to the reader.

In the sections that follow we define composition operators on AD and show that they are precise abstractions of the corresponding operators on NR. We start with the pre- and the post-operator that are rather simple. Then we discuss sequential composition. Afterwards we consider the most interesting and challenging operator: interleaving. Finally, we discuss the abstract semantics of base edges.

4.8 Pre-Operator

We define the (abstract) pre-operator, $pre^\# : AD \rightarrow AD$, as follows:

$$pre^\#(T, D) = \begin{cases} (\emptyset, \emptyset), & \text{if } D = \emptyset \\ (X, \{(\iota, \varphi, \kappa) \in DT \mid (\iota, \varphi, 0) \in D\}), & \text{if } D \neq \emptyset. \end{cases}$$

Lemma 35 *$pre^\#$ is well-defined: for any $(T, D) \in AD$, $pre^\#(T, D) \in AD$.*

PROOF. The only property that is not obvious is that $A := \{(\iota, \varphi, \kappa) \in DT \mid (\iota, \varphi, 0) \in D\}$ is an antichain of short dependence traces. First of all, any dependence trace $(\iota, \varphi, \kappa) \in A$ inherits being short from the dependence trace $(\iota, \varphi, 0) \in D$ that induces its inclusion in A . Secondly, assume that there are distinct dependence traces $\tau, \tau' \in A$ with $\tau \sqsubseteq \tau'$. By the definition of the subsumption order, the transparency bits in τ and τ' must coincide, i.e. we can write them in the form $\tau = (\iota, \varphi, \kappa)$ and $\tau' = (\iota, \varphi', \kappa)$. From $\tau \sqsubseteq \tau'$ it follows that also $(\iota, \varphi, 0) \sqsubseteq (\iota, \varphi', 0)$. But then $(\iota, \varphi, 0)$ and $(\iota, \varphi', 0)$ are two distinct comparable dependence traces in D , which is a contradiction to D being an antichain. Hence $pre^\#(T, D)$ must be an antichain of short dependence traces. \square

The crucial observation for the adequacy of the definition of $pre^\#$ is this.

Lemma 36 *$r \vdash (\iota, \varphi, 0)$ if and only if there is a prefix r' of r with $r' \vdash (\iota, \varphi, \kappa)$.*

PROOF. Let $\varphi = \langle (x_1, y_1), \dots, (x_k, y_k) \rangle$.

‘ \Rightarrow ’: Suppose $r \vdash (\iota, \varphi, 0)$. If $\kappa = 0$, we can choose $r' = r$. So assume $\kappa = 1$.

Choose a decomposition $t_0 \cdot r_1 \cdots r_k \cdot t_k$ of r that witnesses $r \vdash (\iota, \varphi, 0)$. Let $r' = t_0 \cdot r_1 \cdots r_k$. Then, clearly, r' is a prefix of r and $t_0 \cdot r_1 \cdots r_k \cdot t'_k$ with $t'_k = \varepsilon$ is a decomposition of r' that witnesses $r' \vdash (\iota, \varphi, 1)$.

‘ \Leftarrow ’: Suppose r' is a prefix of r with $r' \vdash (\iota, \varphi, \kappa)$. Choose r'' with $r = r' \cdot r''$, and let $t_0 \cdot r_1 \cdots r_k \cdot t_k$ be a decomposition of r' that witnesses $r' \vdash (\iota, \varphi, \kappa)$.

Then $t_0 \cdot r_1 \cdots r_k \cdot t'_k$ with $t'_k = t_k \cdot r''$ is a decomposition of r that witnesses $r \vdash (\iota, \varphi, 0)$. \square

We can now show that the abstract pre-operator is a precise abstraction of the concrete pre-operator.

Theorem 37 (Abstract pre-operator is precise) *Suppose $R \in \text{NR}$. Then $\alpha(\text{pre}(R)) = \text{pre}^\#(\alpha(R))$.*

PROOF. If $R = \emptyset$, then $\alpha(\text{pre}(R)) = \alpha(\emptyset) = (\emptyset, \emptyset) = \text{pre}^\#(\emptyset, \emptyset) = \text{pre}^\#(\alpha(R))$. So let us assume $R \neq \emptyset$.

By unfolding the definitions, we see that $\alpha(\text{pre}(R)) = (T_{\text{pre}(R)}, D_{\text{pre}(R)})$ with

$$\begin{aligned} T_{\text{pre}(R)} &= \{x \mid \exists r, r' \in \text{Runs} : r \cdot r' \in R \wedge r \text{ is transparent for } x\} \\ D_{\text{pre}(R)} &= \{\tau \mid \exists r, r' \in \text{Runs} : r \cdot r' \in R \wedge r \vdash \tau\}^\dagger. \end{aligned}$$

In order to evaluate the right hand side, note first that D_R is non-empty: there is a run $r \in R$ and any such run satisfies $r \vdash (0, \varepsilon, 0)$; moreover, $(0, \varepsilon, 0)$ is \sqsubseteq -maximal and hence contained in D_R . Consequently, the second case applies in the definition of $\text{pre}^\#$ and we have $\text{pre}^\#(\alpha(R)) = \text{pre}^\#(T_R, D_R) = (X, D)$ with

$$D = \{(\iota, \varphi, \kappa) \in \text{DT} \mid (\iota, \varphi, 0) \in D_R\}^\dagger.$$

Thus, we have to show $T_{\text{pre}(R)} = X$ and $D_{\text{pre}(R)} = D$.

$T_{\text{pre}(R)} \subseteq X$ is trivial. In order to see the reverse inclusion, i.e. that $T_{\text{pre}(R)}$ contains any $x \in X$, choose an arbitrary $r \in R$ and observe that the empty run ε is a prefix of r that is transparent for any variable x .

The following chain of implications shows $D_{\text{pre}(R)} \sqsubseteq D$:

$$\begin{aligned} &(\iota, \varphi, \kappa) \in D_{\text{pre}(R)} \\ \Rightarrow & \text{[Equation above, Lemma 19(1.)]} \\ &\exists r, r' \in \text{Runs} : r \cdot r' \in R \wedge r \vdash (\iota, \varphi, \kappa) \\ \text{iff} & \text{[Lemma 36]} \\ &\exists r \in R : r \vdash (\iota, \varphi, 0) \\ \text{iff} & \text{[Set comprehension]} \\ &(\iota, \varphi, 0) \in \{\tau \in \text{DT} \mid \exists r \in R : r \vdash \tau\} \\ \Rightarrow & \text{[Lemma 18, definition } D_R\text{]} \\ &\exists \tau' \in D_R : (\iota, \varphi, 0) \sqsubseteq \tau' \end{aligned}$$

\Rightarrow [See below]
 $\exists \tau \in D : (\iota, \varphi, \kappa) \sqsubseteq \tau$.

The reasoning for the last step is as follows. The subsumption order \sqsubseteq is concerned only with removing gaps from the dependence sequence φ in a dependence trace but leaves the initial and final transparency information untouched. Hence, the dependence trace $\tau' \in D_R$ with $(\iota, \varphi, 0) \sqsubseteq \tau'$ must have the form $\tau' = (\iota, \psi, 0)$. But then $\tau := (\iota, \psi, \kappa) \in D$ and $(\iota, \varphi, \kappa) \sqsubseteq (\iota, \psi, \kappa)$.

Finally, we show $D \sqsubseteq D_{pre(R)}$:

$(\iota, \varphi, \kappa) \in D$
 \Rightarrow [Above equation for D , Lemma 19(1.)]
 $(\iota, \varphi, 0) \in D_R$
 \Rightarrow [Definition of D_R , Lemma 19(1.)]
 $\exists r \in R : r \vdash (\iota, \varphi, 0)$
 iff [Lemma 36]
 $\exists r, r' : r \cdot r' \in R \wedge r \vdash (\iota, \varphi, \kappa)$
 iff [Set comprehension]
 $(\iota, \varphi, \kappa) \in \{\tau \in \mathbf{DT} \mid \exists r, r' : r \cdot r' \in R \wedge r \vdash \tau\}$
 \Rightarrow [Lemma 18]
 $\exists \tau \in \{\tau \in \mathbf{DT} \mid \exists r, r' : r \cdot r' \in R \wedge r \vdash \tau\}^\uparrow : (\iota, \varphi, \kappa) \sqsubseteq \tau$
 iff [Above equation for $D_{pre(R)}$]
 $\exists \tau \in D_{pre(R)} : (\iota, \varphi, \kappa) \sqsubseteq \tau$.

This completes the proof. \square

4.9 Post-Operator

We define the (abstract) post-operator, $post^\# : \mathbf{AD} \rightarrow \mathbf{AD}$, in complete analogy to the pre-operator as follows:

$$post^\#(T, D) = \begin{cases} (\emptyset, \emptyset), & \text{if } D = \emptyset \\ (X, \{(\iota, \varphi, \kappa) \in \mathbf{DT} \mid (\iota, \varphi, \kappa) \in D\}), & \text{if } D \neq \emptyset. \end{cases}$$

By symmetry to the pre-operator we obtain that the post operator is well-defined and a precise abstraction of the post-operator on non-atomic run sets.

Theorem 38 (Abstract post-operator is precise) *Suppose $R \in \mathbf{NR}$. Then $\alpha(post(R)) = post^\#(\alpha(R))$. \square*




Fig. 8. Intuition of sequential composition.

4.10 Sequential Composition

The (*abstract*) *sequential composition operator*, $;\# : \text{AD} \times \text{AD} \rightarrow \text{AD}$, which we write as an infix operator, is defined by

$$(T, D); \# (T', D') = (T \cap T', (D; D')^\uparrow),$$

where

$$D; D' = \{(\iota, \varphi, \kappa) \in D \mid \kappa = 1 \Rightarrow \overrightarrow{\varphi} \in T'\} \quad (1)$$

$$\cup \{(\iota, \varphi, \kappa) \in D' \mid \iota = 1 \Rightarrow \overleftarrow{\varphi} \in T\} \quad (2)$$

$$\cup \{(\iota, \varphi \cdot \psi, \kappa) \in \text{DTS} \mid (\iota, \varphi, 0) \in D, (0, \psi, \kappa) \in D'\} \quad (3)$$

$$\cup \{(\iota, \varphi \cdot \langle(x, z)\rangle \cdot \psi, \kappa) \in \text{DTS} \mid \quad (4)$$

$$\exists y : (\iota, \varphi \cdot \langle(x, y)\rangle, 1) \in D, (1, \langle(y, z)\rangle \cdot \psi, \kappa) \in D'\}.$$

Before we explain the intuition underlying this definition we show well-definedness.

Lemma 39 *The abstract sequential composition operator $;\#$ is well-defined.*

PROOF. We have to show that $(D; D')^\uparrow \in \text{ACS}$ for all $D, D' \in \text{ACS}$, i.e. that $(D; D')^\uparrow$ is an \sqsubseteq -antichain of short dependence traces.

It is easy to see that $D; D'$ (and hence its subset $(D; D')^\uparrow$) contains only short dependence traces: the first two sets contain only dependence traces from D or D' , which consequently are short, and the constructions in the third and fourth set are explicitly restricted to contain short dependence traces. The application of the \uparrow -operator ensures that $(D; D')^\uparrow \in \text{ACS}$ is an \sqsubseteq -antichain. \square

Obviously, a run $r = r' \cdot r''$ composed of two runs r' and r'' is transparent for a variable x if and only if both r' and r'' are. Therefore, transparency information must be intersected in a sequential composition.

Let us explain the intuition underlying the definition of $D; D'$. Suppose given a run $r = r' \cdot r''$ which is composed of two runs $r' \in D$ and $r'' \in D'$ that use distinct virtual variables ($\text{virtual}(r') \cap \text{virtual}(r'') = \emptyset$). Assume that $\tau = (\iota, \varphi, \kappa)$ with $\varphi = \langle d_1, \dots, d_k \rangle$ is a dependence trace of r . Each d_i in φ is a dependence of a sub-piece r_i of r ; we can choose the r_i as short as possible (i.e., such that it starts with an assignment that reads the source variable of d_i and ends with an assignment to the destination variable of d_i). There are four possibilities, how these sub-pieces can be situated in r as illustrated in Fig. 8:

- 1) all of them can lie in r' ;
- 2) all of them can lie in r'' ;
- 3) there is an i , $1 \leq i < k$, such that r_1, \dots, r_i lie in r' and r_{i+1}, \dots, r_k lie in r'' ;
- 4) there is an i such that r_i overlaps with the join point of r' and r'' .

These four cases are handled by the four sets appearing in the definition of $D; D'$:

- 1) in this case, τ is also a dependence trace of r' . Vice versa, dependence traces $\tau' = (\iota', \varphi', \kappa')$ of r' give rise to dependence traces of r . However, if $\kappa' = 1$, no statement that kills $\overrightarrow{\varphi'}$, the destination variable of the last dependence in φ' , is allowed after r_k . Therefore, r' must be transparent for $\overrightarrow{\varphi'}$; hence the side condition in set (1).
- 2) this case is symmetric to case 1).
- 3) in this case, r' has the dependence trace $(\iota, \langle d_1, \dots, d_i \rangle, 0)$ and r'' the dependence trace $(0, \langle d_{i+1}, \dots, d_k \rangle, \kappa)$. Vice versa, dependence traces of r' and r'' of this form give rise to a dependence trace of r .
- 4) choose variables $x, z \in X$ such that $d_i = (x, z)$. Sub-run r_i accomplishes the transfer from x to z via certain intermediate variables. One of these intermediate variables, say y , must bridge the joint point between r' and r'' (i.e., it is assigned to in r' , read from in r'' and not killed in between). As r and r' use distinct virtual variables, y must be a program variable: $y \in X$. Then $(s, \langle d_1, \dots, d_{i-1}, (x, y) \rangle, 1)$ is a dependence trace of r' and $(1, \langle (y, z), d_{i+1}, \dots, d_k \rangle, \kappa)$ is a dependence trace of r'' . The bit 1 as final component of τ' and first component of τ'' is justified, as y is not killed from the place where it is assigned to in r' and read in r'' . Similarly, dependences of r' and r'' of the above form give rise to a dependence trace of r .

It is not hard to see that in all four cases the dependence traces of r' and/or r'' in question are short and \sqsubseteq -maximal if τ is and, vice versa, that each short and \sqsubseteq -maximal dependence trace of r can be composed of short and \sqsubseteq -maximal dependence traces of r' and r'' in the described way.

Lemma 40 (Abstract sequential composition operator is precise)

Suppose $R, S \in \text{NR}$. Then $\alpha(R; S) = \alpha(R); \# \alpha(S)$.

PROOF. By formalizing the intuition described above. \square

4.11 Interleaving

Transparency information for the interleaving $R \otimes S$ of two run sets R and S is easy to obtain from transparency information of the components: a transparent run for a variable x exists in $R \otimes S$ if and only if each component set contains a transparent run. Therefore, the transparency information in T_R and T_S must simply be intersected.

It is far more interesting to consider the dependence traces in $D_{R \otimes S}$ as the two threads modeled by R and S can cooperate in order to exhibit dependences. More specifically, a dependence (u, v) can be composed of complementary dependence sequences of two runs $r \in R$ and $s \in S$, e.g., as illustrated here:

Transfers of r : $u = x_1 \rightarrow y_1 \quad x_2 \rightarrow y_2 \quad x_3 \rightarrow y_4 \quad \cdots \quad x_{k-1} \rightarrow y_{k-1} \quad x_k \rightarrow y_k = v$

Transfers of s : $y_1 \rightarrow x_2 \quad y_2 \rightarrow x_3 \quad \cdots \quad y_{k-1} \rightarrow x_k$

Of course such a combination of complementary dependence sequences can also start and/or end with a dependence of s . And, as a border case, one of the dependence sequences can be empty; the other then just consists of a single dependence. Before we define the abstract interleaving operator, we present in the next section the general definition of when two dependence sequences complement each other to a single dependence and introduce a relation C that extends this definition to dependence traces.

4.11.1 Complementary Dependence Traces

Let $\varphi, \psi \in \text{DS}$ be two dependence sequences (one of them can be empty) and $u, v \in X$. Choose variables such that $\varphi = \langle (x_1, y_1), \dots, (x_k, y_k) \rangle$, $k \geq 0$. We say that ψ *complements* φ to (u, v) if one of the following cases applies:

- (1) $\varphi \neq \varepsilon$, $u = \overleftarrow{\varphi}$, $v = \overrightarrow{\varphi}$, and $\psi = \langle (y_1, x_2), \dots, (y_{k-1}, x_k) \rangle$;
- (2) $\varphi \neq \varepsilon$, $\psi \neq \varepsilon$, $u = \overleftarrow{\varphi}$, $v = \overrightarrow{\psi}$, and $\psi = \langle (y_1, x_2), \dots, (y_{k-1}, x_k), (y_k, v) \rangle$;
- (3) $\varphi \neq \varepsilon$, $\psi \neq \varepsilon$, $u = \overleftarrow{\psi}$, $v = \overrightarrow{\varphi}$, and $\psi = \langle (u, x_1), (y_1, x_2), \dots, (y_{k-1}, x_k) \rangle$; or
- (4) $\psi \neq \varepsilon$, $u = \overleftarrow{\psi}$, $v = \overrightarrow{\psi}$, and $\psi = \langle (u, x_1), (y_1, x_2), \dots, (y_{k-1}, x_k), (y_k, v) \rangle$.




Fig. 9. Complementary dependence traces.

Intuitively, ψ complements φ to (u, v) if the two of them can alternately be combined to a gap-free transfer from u to v . The different cases are distinguished by whether the first read in this gap-free transfer comes from φ (cases 1/2) or ψ (cases 3/4) and whether the last write is in φ (cases 1/3) or ψ (cases 2/4).

Now, consider a dependence trace τ of a run $t \in R \otimes S$ which is an interleaving of the runs $r \in R$, $s \in S$. Then every single dependence in τ must be obtained in the above described fashion from pieces of dependence traces of r and s . We, therefore, generalize this notion of completion to dependence traces as follows. Suppose given dependence traces τ, τ_0, τ_1 , where $\tau = (\iota, \langle (x_1, y_1), \dots, (x_k, y_k) \rangle, \kappa)$, $\tau_0 = (\iota_0, \varphi, \kappa_0)$, $\tau_1 = (\iota_1, \psi, \kappa_1)$. Then we say that τ_1 complements τ_0 to τ , $C(\tau_0, \tau_1, \tau)$ for short, if there are dependence sequences $\varphi_1, \dots, \varphi_k, \psi_1, \dots, \psi_k$ such that

- (1) $\varphi = \varphi_1 \cdot \dots \cdot \varphi_k$ and $\psi = \psi_1 \cdot \dots \cdot \psi_k$;
- (2) ψ_i complements φ_i to (x_i, y_i) for $i = 1, \dots, k$.
- (3) $\iota = 1$ implies $\iota_0 = 1$ and ψ_1 complements φ_1 to (x_1, y_1) according to cases 1 and 2, or $\iota_1 = 1$ and ψ_1 complements φ_1 according to cases 3 and 4; and
- (4) $\kappa = 1$ implies $\kappa_0 = 1$ and ψ_k complements φ_k to (x_k, y_k) according to cases 1 and 3, or $\kappa_1 = 1$ and ψ_k complements φ_k according to cases 2 and 4.

The typical situation of two dependence traces τ_0 and τ_1 that complement each other to a dependence trace τ is illustrated in Fig. 9. For clarity we omit the transparency bits. The dashed vertical lines indicate equality of variables.

A number of elementary properties of the relation C are collected in the following lemma.

Lemma 41 (Basic properties of C) *Suppose $\tau, \tau_0, \tau_1 \in \text{DT}$. Then*

- (1) C is symmetric in the first two parameters: $C(\tau_0, \tau_1, \tau)$ if and only if $C(\tau_1, \tau_0, \tau)$.
- (2) $(0, \varepsilon, 0)$ is a ‘neutral element’: $C((0, \varepsilon, 0), \tau, \tau)$.
- (3) In particular, $C((0, \varepsilon, 0), (0, \varepsilon, 0), (0, \varepsilon, 0))$.

PROOF. Left to the reader. \square

4.11.2 Interleaving Operator

We are now in the position to define the (*abstract*) *interleaving operator*, $\otimes^\# : \text{AD} \times \text{AD} \rightarrow \text{AD}$, which we write again as an infix operator:

$$(T, D) \otimes^\# (T', D') = (T \cap T', \{\tau'' \in \text{DTS} \mid \exists \tau \in D, \tau' \in D' : C(\tau, \tau', \tau'')\}^\uparrow).$$

By restricting the set construction to short dependence traces and applying the $(\cdot)^\uparrow$ operator, the interleaving operator is trivially well-defined. The goal of the remainder of this section is to show that it is a precise abstraction of the interleaving operator on sets of non-atomic runs.

Theorem 42 (Abstract interleaving operator is precise)

Suppose $R, S \in \text{NR}$. Then $\alpha(R \otimes S) = \alpha(R) \otimes^\# \alpha(S)$.

The proof is deferred to Section 4.11.5. Before that, we establish a number of lemmas that capture the main insights underlying the proof.

4.11.3 Soundness Lemmas

The lemmas in this section are concerned with the soundness of the abstract interleaving composition operator, i.e. they are crucial for the proof that $\alpha(R \otimes S) \sqsubseteq \alpha(R) \otimes^\# \alpha(S)$ for any two run sets R, S . The critical point here is to guarantee that our definition of the abstract interleaving operator includes enough dependence traces.

As a first step, we show that each dependence trace of some interleaving of two runs r, s can also be obtained by combining two dependence traces of the component runs r and s via the relation C .

Let $r, s, t \in \text{Runs}$ with $\text{virtual}(r) \cap \text{virtual}(s) = \emptyset$ and $\tau \in \text{DT}$.

Lemma 43 *Suppose $t \in r \otimes s$ and $t \vdash \tau$. Then there are $\tau_r, \tau_s \in \text{DT}$ with $r \vdash \tau_r$, $s \vdash \tau_s$, and $C(\tau_r, \tau_s, \tau)$.*

PROOF. Assume that t is an interleaving of r and s and $\tau = (\iota, \langle d_1, \dots, d_k \rangle, \kappa)$ is a dependence trace of t . Each d_i is a dependence of a certain sub-run t_i of t and each t_i is an interleaving of certain sub-runs of r and s .

From t_i we can construct dependence traces φ_i and ψ_i of these sub-runs of r and s such that φ_i complements ψ_i to dependence d_i . This is described below. Then $\varphi_1 \cdot \dots \cdot \varphi_k$ and $\psi_1 \cdot \dots \cdot \psi_k$ are dependence sequences of r and s , resp., and we can choose transparency bits $\iota_r, \kappa_r, \iota_s, \kappa_s \in \mathbb{B}$ such that $\tau_r = (\iota_r, \varphi_1 \cdot \dots \cdot \varphi_k, \kappa_r)$ and $\tau_s = (\iota_s, \psi_1 \cdot \dots \cdot \psi_k, \kappa_s)$ are dependence traces of r and s , resp., such that $C(\tau_r, \tau_s, \tau)$ holds. Specifically, we choose $\iota_r = \iota$ if the first




Fig. 10. Dependence traces of interleavings are induced by complementary dependence traces of the components.

assignment instance involved in the mediation of d_1 belongs to r and $\iota_s = \iota$ if it belongs to s , and similarly for the final transparency bits and the last assignment instance involved in the mediation of d_k . All other transparency bits are chosen 0.

Let us now explain how to construct the dependence sequences φ_i and ψ_i mentioned above. Choose program variables x, y such that $d_i = (x, y)$. Sub-run t_i of t exhibits d_i via certain assignment instances $a_j := e_j$, $j = 1, \dots, l$. In particular, $a_l = y$. Each of these assignment instances lies either in a sub-piece of r or a sub-piece of s . Let us consider the case that the first assignment instance $a_1 := e_1$ lies in a sub-piece of r ; the case that it lies in a sub-piece of s is analogous. We can then find indices $0 < j_0 < j_1 < \dots < j_n$ such that $a_j := e_j$ lies in a sub-piece of r if $j_m < j \leq j_{m+1}$ for an *even* $m \in \{0, \dots, n-1\}$ and in a sub-piece of s otherwise. In particular, for any $j \in \{j_1, \dots, j_{n-1}\}$ one of the assignments instances $a_j := e_j$ and $a_{j+1} := e_{j+1}$ lies in a sub-piece of r and the other one in a sub-piece of s . This implies that a_j must be a program variable, because it appears in e_{j+1} and $\text{virtual}(r) \cap \text{virtual}(s) = \emptyset$. Choose now

$$\begin{aligned}\varphi_i &= \langle (x, a_{j_1}), (a_{j_2}, a_{j_3}), \dots, (a_{j_{n-2}}, a_{j_{n-1}}) \rangle, \\ \psi_i &= \langle (a_{j_1}, a_{j_2}), (a_{j_3}, a_{j_4}), \dots, (a_{j_{n-1}}, y) \rangle\end{aligned}$$

if n is even and

$$\begin{aligned}\varphi_i &= \langle (x, a_{j_1}), (a_{j_2}, a_{j_3}), \dots, (a_{j_{n-1}}, y) \rangle, \\ \psi_i &= \langle (a_{j_1}, a_{j_2}), (a_{j_3}, a_{j_4}), \dots, (a_{j_{n-2}}, a_{j_{n-1}}) \rangle\end{aligned}$$

if n is odd. Then φ_i and ψ_i are dependence sequences of the sub-runs of r and s that comprise t_i and, obviously, φ_i complements ψ_i to d_i . \square

Example 44 Fig. 10 illustrates the construction in the proof of Lemma 43. The run t is an interleaving of the runs r and s . We can thus decompose r

and s into sub-runs such that t is obtained by alternately shuffling these sub-runs together; in the example $r = r_1 \cdot r_2 \cdot r_3 \cdot r_4 \cdot r_5$, $s = s_1 \cdot s_2 \cdot s_3 \cdot s_4$, and $t = r_1 \cdot s_1 \cdot r_2 \cdot s_2 \cdot r_3 \cdot s_3 \cdot r_4 \cdot s_4 \cdot r_5$.

Let us assume that $\tau = (\iota, \langle (a, b), (c, d) \rangle, \kappa)$ is a dependence trace of t . Then there are sub-runs t_1 and t_2 of t that exhibit the two dependences (a, b) and (c, d) , e.g., as shown in the figure. These sub-runs overlap in a certain way with the decompositions of r and s ; in the example in the figure, for instance, t_1 overlaps with a postfix of r_1 , all of s_1, r_2, s_2 , and a prefix of r_3 . The dependence (a, b) is exhibited via certain intermediate assignments $a_i := e_i$ (not shown in the figure); we call these assignments *crucial* in the following.

There may be sub-runs of r and/or s that overlap with t_i but do not contain a crucial assignment. Such sub-runs must be transparent for the variable that transfers the dependence at this moment and can be ignored. In our example, r_2 is such a sub-run and g is the variable that transfers the dependence while r_2 is executed.

Whenever two successive crucial assignments lie in sub-pieces of different runs, the dependence must be transferred in a program variable between these assignments because r and s do not share virtual variables. In the figure, e.g., e is the variable that transfers the dependence from the last crucial assignment in r_1 to the first crucial assignment in s_1 and f transfers it from the last crucial assignment in s_1 to the first crucial assignment in r_2 . From these variables we can construct dependence traces τ_r and τ_s of r and s such that $C(\tau_r, \tau_s, \tau)$ holds. In Fig. 10, for instance, we have $\tau_r = (\iota, \langle (a, e), (f, b), (h, i), (j, d) \rangle, \kappa)$ and $\tau_s = (0, \langle (e, f), (c, h), (i, j) \rangle, 0)$. \square

Lemma 43 ensures that combining dependence traces of component runs via C is fundamentally rich enough to give us all dependence traces of potential interleavings. However, in our abstract domain, we do not collect *all* dependence traces but only the *maximal* ones. Therefore, we only combine the maximal dependence traces of component runs in the definition of interleaving, which is the best we can do with the available information. Can we really obtain all the *maximal* dependence traces just from the *maximal* dependence traces of the components?

The next lemma provides us with a kind of shortening rule that is crucial for the proof that maximal dependence traces of component run sets suffice to infer the maximal dependence traces of their interleaving.

Suppose $\tau_0, \tau'_0, \tau_1, \tau \in \text{DT}$.

Lemma 45 *Suppose $C(\tau_0, \tau_1, \tau)$ and $\tau_0 \sqsubseteq \tau'_0$. Then there are dependence traces $\tau'_1, \tau' \in \text{DT}$ such that $\tau_1 \leq \tau'_1$, $\tau \sqsubseteq \tau'$, and $C(\tau'_0, \tau'_1, \tau')$. By symmetry of C (Lemma 41(1)) an analogous property holds with the roles of τ_0 and*

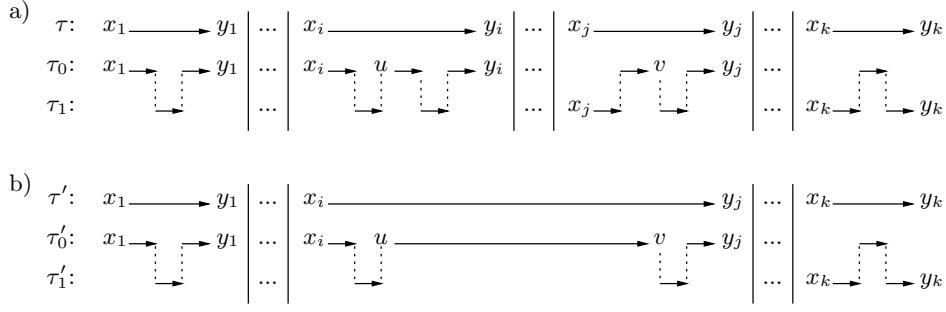


Fig. 11. Removing gaps in a component dependence trace

τ_1 exchanged.

PROOF. The proof is illustrated in Fig. 11. In a) the typical situation of dependence traces τ_0, τ_1 and τ with $C(\tau_0, \tau_1, \tau)$ is shown. For clarity the transparency bits are omitted. In b) a typical dependence trace τ'_0 with $\tau_0 \sqsubseteq \tau'_0$ is shown. It is obtained from τ_0 by removing all gaps between the target variable u of a certain dependence d in τ_0 and the destination variable v of a later dependence e . We can remove all the dependences from τ_1 that are used to fill some or all of these gaps in $C(\tau_0, \tau_1, \tau)$. This results in a dependence trace τ'_1 with $\tau_1 \leq \tau'_1$ as shown in b). Then the dependence traces τ'_0 and τ'_1 complement each other to a dependence trace τ' with $\tau \leq \tau'$ as shown. As border cases, we may have $\tau'_0 = \tau_0$, if none of the gaps between d and e is filled in $C(\tau_0, \tau_1, \tau)$, or $\tau' = \tau$ if d and e are used in $C(\tau_0, \tau_1, \tau)$ in the same dependence of τ . But this does not invalidate our reasoning as \sqsubseteq and \leq are reflexive. \square

By iteratively applying this shortening rule, we obtain the following lemma that is of direct use in the proof of Theorem 42.

Lemma 46 *Suppose $\tau_r \in \{\tau \mid \exists r \in R : r \vdash \tau\}$, $\tau_s \in \{\tau \mid \exists s \in S : s \vdash \tau\}$, and $C(\tau_r, \tau_s, \tau)$. Then there are $\tau'_r \in D_R$, $\tau'_s \in D_S$, and $\tau' \in \text{DT}$ with $C(\tau'_r, \tau'_s, \tau')$ and $\tau \sqsubseteq \tau'$.*

PROOF. The problem is that τ_r and τ_s need not be \sqsubseteq -maximal in their respective set. Hence they may not belong to D_R and D_S , respectively. By iteratively applying Lemma 45, however, we can determine dependence traces τ_r^\uparrow and τ_s^\uparrow that are \sqsubseteq -maximal in these sets (and hence belong to D_R and D_S , respectively) as well as a dependence trace τ^\uparrow with $C(\tau_r^\uparrow, \tau_s^\uparrow, \tau^\uparrow)$ and $\tau \sqsubseteq \tau^\uparrow$:

We start with $(\tau_r^\uparrow, \tau_s^\uparrow, \tau^\uparrow) := (\tau_r, \tau_s, \tau)$. This initialization trivially ensures $\tau_r^\uparrow \in \{\tau \mid \exists r \in R : r \vdash \tau\}$, $\tau_s^\uparrow \in \{\tau \mid \exists s \in S : s \vdash \tau\}$, $C(\tau_r^\uparrow, \tau_s^\uparrow, \tau^\uparrow)$, and $\tau \sqsubseteq \tau^\uparrow$, which is an invariant of the loop we describe in the following.

If τ_r^\uparrow is not \sqsubseteq -maximal in $\{\tau \mid \exists r \in R : r \vdash \tau\}$, we can choose a dependence trace $\tau'_r \in \{\tau \mid \exists r \in R : r \vdash \tau\}$ which is strictly larger: $\tau_r^\uparrow \sqsubset \tau'_r$. Then, by Lemma 45, there are τ'_s and τ' with $\tau_s^\uparrow \leq \tau'_s$, $\tau \sqsubseteq \tau^\uparrow \sqsubseteq \tau'$, and $C(\tau'_r, \tau'_s, \tau')$. By Proposition 14, $\tau'_s \in \{\tau \mid \exists r \in R : r \vdash \tau\}$, hence the invariant remains valid. We then set $(\tau_r^\uparrow, \tau_s^\uparrow, \tau^\uparrow) := (\tau'_r, \tau'_s, \tau')$. We can proceed analogously, if τ_s^\uparrow is not maximal in $\{\tau \mid \exists s \in S : s \vdash \tau\}$.

This shortening procedure is applied iteratively until both τ_r^\uparrow and τ_s^\uparrow are \sqsubseteq -maximal in their respective sets. Termination is guaranteed, because in each step either the dependence sequence in τ_r^\uparrow or in τ_s^\uparrow becomes shorter and the dependence sequence in the other dependence trace does not become longer. \square

4.11.4 Completeness Lemmas

The lemmas in this section are concerned with completeness of the interleaving operator, i.e. they are important for the proof that $\alpha(R \otimes S) \sqsupseteq \alpha(R) \otimes^\# \alpha(S)$ for any two non-atomic run sets R, S . They crucially depend on runs being non-atomic.

A dependence of a non-atomic run r must involve a virtual variable at a certain stage as assignments that have program variables on both the left- and the right-hand-side do not occur in non-atomic runs. But when the execution of r is in such a stage, no parallel thread can disturb propagation of the dependence because parallel threads do not interfere on virtual variables. This observation underlies the proof of the following lemma.

Lemma 47 *Suppose r, s are runs with $\text{virtual}(r) \cap \text{virtual}(s) = \emptyset$, and $x, y \in X$. If r exhibits (x, y) then there is a run $t \in r \otimes s$ that exhibits (x, y) .*

PROOF. Suppose (x, y) is a dependence of r . This means that r can be written in the form $r = r_0 \cdot \langle a_1 := e_1 \rangle \cdot r_1 \cdot \langle a_2 := e_2 \rangle \cdot r_2 \cdot \dots \cdot r_{l-1} \cdot \langle a_l := e_l \rangle \cdot r_l$ as in the definition of “ r exhibits (x, y) ”. Then in particular e_1 contains the variable x . As x is a program variable, this implies by the form of assignments appearing in runs that a_1 must be a virtual variable (cf. the definition of **Asg**). As $\text{virtual}(r) \cap \text{virtual}(s) = \emptyset$, s must thus be transparent for a_1 . Hence the run $t \in r \otimes s$ defined by

$$t := r_0 \cdot \langle a_1 := e_1 \rangle \cdot s \cdot r_1 \cdot \langle a_2 := e_2 \rangle \cdot r_2 \cdot \dots \cdot r_{l-1} \cdot \langle a_l := e_l \rangle \cdot r_l$$

still exhibits dependence (x, y) . \square

Note that this argument crucially depends on the assumption about the form of assignments in runs that derives from the assumption that assignments

execute non-atomically. If assignments execute atomically, the above lemma is no longer valid.

Example 48 Consider the parallel execution of the two straight-line programs $\pi_1 = (y := x)$ and $\pi_2 = (x := 0; y := 0)$.

If assignment statements execute atomically, there are just three possible runs,

- 1) $\langle x := 0, y := 0, y := x \rangle$,
- 2) $\langle x := 0, y := x, y := 0 \rangle$, and
- 3) $\langle y := x, x := 0, y := 0 \rangle$.

None of these runs exhibits dependence (x, y) because either x is killed before $y := x$ is executed as in 1) and 2), or y is killed after $y := x$ is executed as in 2) and 3).

If, on the other hand, assignment statements may execute non-atomically, then the two initialization statements in π_2 could well be executed after x is read but before y is written. This is witnessed by the run

- 4) $\langle v := x, x := 0, y := 0, y := v \rangle$,

where v is a virtual variable, in our model of non-atomic execution. In contrast to the runs 1)-3), run 4) exhibits dependence (x, y) . \square

Lemma 47 provides an intuitive explanation why precise analysis of parallel programs is simpler if we assume non-atomic execution of assignments. Under this assumption dependences once generated by a thread cannot be definitely destroyed by its environment. Thus, an analysis that collects positive information about potential dependences is precise. (In order to do this in a \downarrow compositional fashion it must collect more information, namely (maximal, short) dependence traces. If we analyze with respect to the assumption that assignments execute atomically, there is a complex interplay between the way dependences are generated by a thread and the order of re-initializations performed by its environment as illustrated by the above example. Therefore, an analysis that just collects positive information is doomed to be imprecise.

Lemma 49 Suppose r_0, r_1 are runs with $\mathit{virtual}(r_0) \cap \mathit{virtual}(r_1) = \emptyset$ and τ_0, τ_1, τ are dependence traces with $r_0 \vdash \tau_0$, $r_1 \vdash \tau_1$, and $C(\tau_0, \tau_1, \tau)$. Then there is a run $r \in r_0 \otimes r_1$ such that $r \vdash \tau$.

PROOF. For notational convenience, we discuss the case that the dependence sequence in τ consists of just a single dependence; the generalization to arbitrary dependence sequences is left to the reader. Let $\tau = (\iota, \langle (u, v) \rangle, \kappa)$. Furthermore, let $\tau_0 = (\iota_0, \phi, \kappa_0)$ and $\tau_1 = (\iota_1, \psi, \kappa_1)$.

Let us assume that case 2 in the definition of $C(\tau_0, \tau_1, \tau)$ applies; the other cases are similar. Then we can choose variables $u = x_1, \dots, x_{k+1} = v$ such that

$$\varphi = \langle (x_1, y_1), \dots, (x_k, y_k) \rangle \quad \text{and} \quad \psi = \langle (y_1, x_2), \dots, (y_k, x_{k+1}) \rangle,$$

and it is $\iota_0 = 1$ if $\iota = 1$ and $\kappa_1 = 1$ if $\kappa = 1$. As $r_0 \vdash \tau_0$ and $r_1 \vdash \tau_1$ we can write r_0 and r_1 in the form

$$r_0 = t_0^0 \cdot r_1^0 \cdot t_1^0 \cdot r_2^0 \cdots t_{k-1}^0 \cdot r_k^0 \cdot t_k^0 \quad \text{and} \quad r_1 = t_0^1 \cdot r_1^1 \cdot t_1^1 \cdot r_2^1 \cdots t_{k-1}^1 \cdot r_k^1 \cdot t_k^1$$

such that

- 1) r_i^0 exhibits (x_i, y_i) and r_i^1 exhibits (y_i, x_{i+1}) for $i = 1, \dots, k$;
- 2) t_0^0 is transparent for u if $\iota = 1$ (and hence $\iota_0 = 1$); and
- 3) t_k^1 is transparent for v if $\kappa = 1$ (and hence $\kappa_1 = 1$).

The run $r_1^0 \cdot r_1^1 \cdot r_2^0 \cdot r_2^1 \cdots r_k^0 \cdot r_k^1$ clearly exhibits dependence (u, v) , but in order to construct an interleaving of r_0 and r_1 , we must also execute the intermediate code pieces t_i^j . Fortunately, each of the dependences realized by some r_i^j must involve a virtual variable; and, while the transfer is in such a stage, code pieces of the other run, r_{1-j} , can safely be executed without destroying the dependence, due to the disjointness of the virtual variables used in r_0 and r_1 . Thus, we can execute each code piece t_i^1 at such a stage of execution of r_{i+1}^0 and, similarly, t_i^0 during such a stage of r_i^1 . The rest of the proof pursues this argument more formally.

By Lemma 47, there are interleavings $s_i^0 \in r_i^0 \otimes t_{i-1}^1$ and $s_i^1 \in r_i^1 \otimes t_i^0$ such that, for $i = 1, \dots, k$, s_i^0 still exhibits (x_i, y_i) and s_i^1 still exhibits (y_i, x_{i+1}) . Then the run $r := t_0^0 \cdot s_1^0 \cdot s_1^1 \cdot s_2^0 \cdot s_2^1 \cdots s_k^0 \cdot s_k^1 \cdot t_k^1$ is an interleaving of r_0 and r_1 (i.e. $r \in r_1 \otimes r_2$). On the other hand, $r \vdash \tau$ because $s_1^0 \cdot s_1^1 \cdot s_2^0 \cdot s_2^1 \cdots s_k^0 \cdot s_k^1$ exhibits dependence (u, v) and items 2) and 3) above give the transparency properties. \square

Like Lemma 47, Lemma 49 fails to hold if assignments execute atomically as illustrated by the following example.

Example 50 Consider the two programs $\pi_1 = (y := x)$ and $\pi_2 = (x := 0; y := 0; z := y)$ and the three dependence traces $\tau_1 = (1, \langle (x, y) \rangle, 1)$, $\tau_2 = (1, \langle (y, z) \rangle, 1)$, and $\tau = (1, \langle (x, z) \rangle, 1)$.

If assignments execute atomically, π_1 has only the run $r_1 = \langle y := x \rangle$ and π_2 has only the run $r_2 = \langle x := 0, y := 0, z := y \rangle$. Clearly, τ_1 is a dependence trace of r_1 and τ_2 is a dependence trace of r_2 , independently of whether assignments execute atomically or not. Moreover, $C(\tau_1, \tau_2, \tau)$ holds.

But only the following four runs are possible interleavings of r_1 and r_2 :

- 1) $\langle x := 0, y := 0, z := y, y := x \rangle$,
- 2) $\langle x := 0, y := 0, y := x, z := y \rangle$,
- 3) $\langle x := 0, y := x, y := 0, z := y \rangle$, and
- 4) $\langle y := x, x := 0, y := 0, z := y \rangle$.

It is not hard to see that τ is not exhibited by any of these runs.

If, on the other hand, assignments do not execute atomically, there are also runs like

- 5) $\langle v := x, x := 0, y := 0, y := v, u := y, z := u \rangle$,

where u, v are virtual variables, which exhibits dependence trace τ . \square

4.11.5 Proof of Theorem 42

We can now put the pieces together and prove Theorem 42. By unfolding the definitions, we have

$$\begin{aligned}\alpha(R \otimes S) &= (T_{R \otimes S}, D_{R \otimes S}) \text{ and} \\ \alpha(R) \otimes^\# \alpha(S) &= (T_R \cap T_S, D),\end{aligned}$$

where $D = \{\tau \in \text{DTS} \mid \exists \tau_R \in D_R, \tau_S \in D_S : C(\tau_R, \tau_S, \tau)\}^\uparrow$. Consequently, we have to show $T_{R \otimes S} = T_R \cap T_S$ and $D_{R \otimes S} = D$.

“ $T_{R \otimes S} \subseteq T_R \cap T_S$ ”: If $x \in T_{R \otimes S}$, then there is a run $t \in R \otimes S$ that is transparent for x . By definition, t is an interleaving of runs $r \in R$ and $s \in S$.

These runs r, s must then also be transparent for x . Thus, $x \in T_R \cap T_S$.

“ $T_{R \otimes S} \supseteq T_R \cap T_S$ ”: If $x \in T_R \cap T_S$, then there are runs $r \in R$ and $s \in S$ that are transparent for x . By bounded renaming of virtual variables these runs can be chosen such that they do not share virtual variables. Then all interleavings of these two runs are in $S \otimes R$, and all of them are transparent for x . Thus, $x \in T_{R \otimes S}$.

“ $D_{R \otimes S} \sqsubseteq D$ ”: In order to show this relationship, assume that we are given $\tau \in D_{R \otimes S}$. Then we have, by the definition of $D_{R \otimes S}$ and Lemma 19(1.):

$$\begin{aligned}& \exists t \in R \otimes S : t \vdash \tau \\ \text{iff} & \quad [\text{Definition } R \otimes S] \\ & \exists r \in R, s \in S, t \in r \otimes s : t \vdash \tau \\ \Rightarrow & \quad [\text{Lemma 43}] \\ & \exists r \in R, s \in S, \tau_r, \tau_s \in \text{DT} : r \vdash \tau_r \wedge s \vdash \tau_s \wedge C(\tau_r, \tau_s, \tau)\end{aligned}$$

\Rightarrow [Shunting, set comprehension]
 $\exists \tau_r \in \{\tau \mid \exists r \in R : r \vdash \tau\}, \tau_s \in \{\tau \mid \exists s \in S : s \vdash \tau\} : C(\tau_r, \tau_s, \tau)$.
 \Rightarrow [Lemma 46]
 $\exists \tau_r \in D_R, \tau_s \in D_S, \tau' \in \text{DT} : C(\tau_r, \tau_s, \tau') \wedge \tau \sqsubseteq \tau'$
 iff [Set comprehension, see below]
 $\exists \tau' \in \{\tau \in \text{DTS} \mid \exists \tau_R \in D_R, \tau_S \in D_S : C(\tau_R, \tau_S, \tau)\} : \tau \sqsubseteq \tau'$
 \Rightarrow [Definition D , Lemma 18]
 $\exists \tau' \in D : \tau \sqsubseteq \tau'$.

In the step marked “see below”, we must prove for “ \Rightarrow ” that τ' can be chosen as a *short* dependence trace, which is *not* true for this step in isolation. But, it is true under the assumption that $\tau \in D_{R \otimes S}$ which underlies the whole calculation: as a consequence of this assumption τ is *short* and this implies that any τ' with $\tau \sqsubseteq \tau'$ must also be short (Lemma 27). A calculation, in which this step is valid in isolation, requires to furnish each of the preceding predicates with the conjunct $\tau \in D_{R \otimes S}$, which would clutter the calculation.

“ $D_{R \otimes S} \sqsupseteq D$ ”: This is shown by the following chain of implications:

$\tau \in D$
 \Rightarrow [Definition of D , Lemma 19(1.)]
 $\exists \tau_R \in D_R, \tau_S \in D_S : C(\tau_R, \tau_S, \tau)$
 \Rightarrow [Definition D_R, D_S , Lemma 19(1.)]
 $\exists r \in R, s \in S, \tau_R, \tau_S : r \vdash \tau_R \wedge s \vdash \tau_S \wedge C(\tau_R, \tau_S, \tau)$
 iff [By bounded renaming of virtual variables in s]
 $\exists r \in R, s \in S, \tau_R, \tau_S :$
 $r \vdash \tau_R \wedge s \vdash \tau_S \wedge C(\tau_R, \tau_S, \tau) \wedge \text{virtual}(r) \cap \text{virtual}(s) = \emptyset$
 \Rightarrow [Lemma 49, definition $R \otimes S$]
 $\exists t \in R \otimes S : t \vdash \tau$
 iff [Set comprehension]
 $\tau \in \{\tau \in \text{DT} \mid \exists r \in R \otimes S : r \vdash \tau\}$
 \Rightarrow [Lemma 18, definition $D_{R \otimes S}$]
 $\exists \tau' \in D_{R \otimes S} : \tau \sqsubseteq \tau'$.

This ends the proof of Theorem 42. \square

4.12 Base Edges

In Section 3 we discussed that the atomicity assumptions about assignments may vary and that this gives rise to different definitions of the non-atomic run sets $\llbracket x := e \rrbracket$ assigned to an assignment statement $x := e$. Fortunately, all reasonable choices result in the same abstraction which is given by the following definition:

$$\llbracket x := e \rrbracket^\# = (X \setminus \{x\}, \{(\iota, \langle (y, x) \rangle, \kappa) \mid \iota, \kappa \in \mathbb{B}, y \text{ appears in } e\}).$$

Whatever atomicity assumption we are working with, all runs in $\llbracket x := e \rrbracket$ will contain certain auxiliary assignments to virtual variables and a single assignment to x . No program variable except x will ever be the target of an assignment in a run in $\llbracket x := e \rrbracket$. Hence, all non-atomic runs are transparent just for the program variables in $X \setminus \{x\}$, which explains the adequacy of the first component of $\llbracket x := e \rrbracket^\#$. Moreover, it implies that no dependence trace of a non-atomic run can embody a dependence sequence that is longer than one or has a destination variable different from x . Each reasonable non-atomic run induces the same dependences between program variables as $x := e$, hence the induced dependences are (y, x) where y is a variable appearing in e . Moreover, no reasonable run kills a variable in e before it reads it or kills x after it has written it, which implies that the transparency bits can be chosen arbitrarily.

All dependence traces included in the second component of $\llbracket x := e \rrbracket^\#$ are trivially short and \sqsubseteq -maximal, which implies well-definedness.

Proposition 51 *Suppose $x := e \in \text{Stmt}$. Then $\alpha(\llbracket x := e \rrbracket) = \llbracket x := e \rrbracket^\#$. \square*

Statement **skip** has just the single run ε , which is obviously transparent for all variables and has just the dependence trace $(0, \varepsilon, 0)$. Hence, we define $\llbracket \text{skip} \rrbracket^\# = (X, \{(0, \varepsilon, 0)\})$.

Proposition 52 $\alpha(\llbracket \text{skip} \rrbracket) = \llbracket \text{skip} \rrbracket^\#$. \square

We define the abstract interpretation of a base edge e of the underlying flow graph as the interpretation of the statement $A(e)$ associated with e : $\llbracket e \rrbracket^\# = \llbracket A(e) \rrbracket^\#$.

Proposition 53 $\alpha(\llbracket e \rrbracket) = \llbracket e \rrbracket^\#$ for all base edges e . \square

4.13 Run-Time

In this section we show that we can compute the abstract operations $pre^\#$, $post^\#$, $;\#$, and $\otimes^\#$ in time $2^{p(|X|)}$, where $p(x)$ is a polynomial. We emphasize

that we do *neither* intend to develop efficient implementations of the operations *nor* to present a very precise analysis. The results of this section will mainly be used in order to establish the qualitative complexity statement that the algorithms developed later run in exponential time. We are, however, interested in uncovering the parameter of exponential growth: it is the number of program variables $|X|$ rather than the size of the parallel flow graph.

Let us investigate the most expensive operation, interleaving, to some detail. First of all, we recall its definition from Section 4.11:

$$(T, D) \otimes^\# (T', D') = (T \cap T', D''^\uparrow),$$

where $D'' = \{\tau'' \in \text{DTS} \mid \exists \tau \in D, \tau' \in D' : C(\tau, \tau', \tau'')\}$. The sets T and T' are subsets of X , the set of program variables. Computing the intersection of T and T' is cheap: if we represent these sets as bit-strings (of length $|X|$), we can clearly calculate the intersection in time $\mathcal{O}(|X|)$ by looking through the bit-strings for T and T' once.

D and D' are antichains of short dependence traces, hence $D, D' \subseteq \text{DTS}$. By Lemma 25, the cardinality of DTS and hence of D and D' is $\mathcal{O}(|X|^{2|X|+2})$. This clearly is $\mathcal{O}(2^{p_0(|X|)})$ for some polynomial $p_0(x)$. We can hence consider at most $\mathcal{O}(2^{2p_0(|X|)})$ pairs of dependence traces τ and τ' when computing D'' . For each fixed pair of dependence traces τ, τ' all dependence traces τ'' with $C(\tau, \tau', \tau'')$ can be determined in time $\mathcal{O}(2^{p_1(|X|)})$ for some polynomial $p_1(x)$. We leave it to the reader to invent some procedure for this task that realizes this rather brutal bound. Even a very naive procedure that lists all short dependence traces τ'' and then checks for each listed dependence trace whether $C(\tau, \tau', \tau'')$ holds will do. The observation that τ, τ' , and τ'' are short, and hence the length of their dependence sequences is bounded by $|X| + 1$ is helpful. As a consequence, we can calculate D'' in time $\mathcal{O}(2^{2p_0(|X|)+p_1(|X|)})$. Again $\mathcal{O}(2^{p_0(|X|)})$ is an asymptotic bound for the size of D'' because $D'' \subseteq \text{DTS}$. It is, therefore, not hard to see that D''^\uparrow , the second component of $(T, D) \otimes^\# (T', D')$, can be computed from D'' in time $\mathcal{O}(2^{p_2(|X|)})$ for some polynomial $p_2(x)$. Hence the overall cost of computing $(T, D) \otimes^\# (T', D')$ is $\mathcal{O}(2^{p(|X|)})$ for some polynomial $p(x)$.

By similar considerations we can show that the other operations can be computed in time $\mathcal{O}(2^{p(|X|)})$ too.

Lemma 54 *The operations $\text{pre}^\#, \text{post}^\#, ;^\#,$ and $\otimes^\#$ can be computed in time $\mathcal{O}(2^{p(|X|)})$ for some polynomial $p(x)$. \square*




Fig. 12. The situation in the transfer lemma

4.14 Discussion

In this section, we have defined an abstraction of sets of non-atomic runs from which the exhibited dependences can be derived. Run sets are abstracted to antichains of short dependence traces that capture the potential to exhibit dependences in cooperation with a parallel environment. The abstraction also records the set of program variables for which a transparent run exists in the abstracted run set. This information is necessary to propagate the transparency bits of the dependence traces properly in sequential contexts. We have defined abstract interpretations of the operations and constants used in the constraint systems of Section 2.5 and have shown that they precisely abstract the corresponding operations on sets of non-atomic runs.

5 Detecting Copy Constants and Eliminating Faint Code

In this section we show that we can detect copy constants and eliminate faint code in parallel flow graphs completely relative to the non-atomic semantics. The basic idea is to evaluate the constraint system for bridging runs over the abstract domain \mathbf{AD} from the previous section and to exploit this information.

We have seen that the abstract counterparts of the operators and constants appearing in the constraint systems in Section 2 abstract the corresponding operators on non-atomic run sets precisely. Moreover, the abstraction mapping $\alpha : \mathbf{NR} \rightarrow \mathbf{AD}$ is universally disjunctive (Proposition 33). This implies that the least solution of the constraint systems over domain \mathbf{AD} consists just of the abstractions of the least solution over domain \mathbf{NR} . This is commonly known in the area of abstract interpretation [21,22] and follows directly from the following fixpoint-theoretic lemma known as Transfer-Lemma [29] or μ -Fusion Rule [30].

Lemma 55 (Transfer lemma) *Suppose L, L' are complete lattices, $f : L \rightarrow L$ and $g : L' \rightarrow L'$ are monotonic functions and $h : L \rightarrow L'$ (Fig. 12).*

If h is universally disjunctive and $h \circ f = g \circ h$ then $h(\mu f) = \mu g$, where μf and μg are the least fixpoints of f and g , respectively. \square

The least solution of a constraint system over some domain corresponds in a straightforward way to the least fixpoint of a function derived from the

constraints. The facts recalled above ensure that the premises of the Transfer-Lemma hold for the functions f and g derived from the concrete and abstract interpretation of the constraint systems over non-atomic runs and over \mathbf{AD} , respectively, and the transfer function h that component-wise maps the concrete interpretation x of each variable X of the constraint system to its abstraction $\alpha(x)$. As \mathbf{AD} is finite, we can compute the least solution of the constraint system for (non-atomic) bridging runs over lattice \mathbf{AD} effectively by fixpoint iteration. From the computed values we can read off in particular all the dependences of the bridging runs: if $\alpha(R) = (T, D)$ is the precise abstraction of a set R of (non-atomic) runs then (x, y) is a dependence of a run in R if and only if $(1, \langle(x, y)\rangle, 1) \in D$ (Proposition 12).

Based on this information we can detect copy constants and eliminate faint code. The corresponding algorithms run in exponential time. Indeed the point here is *not* to develop efficient algorithms—we will see in the next section that all these problems are intractable already for loop-free parallel programs—the point is that these problems can be solved effectively at all! This comes as a surprise, because the corresponding problems are uncomputable, if we assume atomic execution of assignments [12].

Without further ado, we present, in the remainder of this section, the algorithms for detection of copy constants (Section 5.1) and faint code elimination (Section 5.2). After the presentation of the algorithms, we analyze their asymptotic run-time in Section 5.3 and finish the section with some concluding remarks.

5.1 Copy Constant Detection

A variable x is a copy-constant at a program point u if it gets assigned the same value on all runs reaching u either through a constant assignment (like in $\langle x := 42 \rangle$) or a constant assignment followed by copying assignments (like in $\langle z := 42, y := z, x := y \rangle$). Of course the runs may contain other assignments also that do not influence the final value of x (like in $\langle x := 42, y := a + b \rangle$). Thus, in copy constant detection only assignments of the simple form $x := k$, where k is a constant or variable, are interpreted, all other forms of assignments (e.g. $x := y + 1$) are (conservatively) assumed to make x non-constant [31].

Algorithm 1 in Fig. 13 reads a parallel flow graph, a program point $v \in N$, and a program variable $y \in X$ and decides whether y is a copy constant at v or not. For this purpose it first computes (in Steps 1 and 2) for each program point w the set

$$I[w] = \{x \mid e_{Main} \Longrightarrow c_w \xrightarrow{r} c_v, At_w(c_w), At_v(c_v), \hat{r} \text{ exhibits dep. } (x, y)\}.$$

Algorithm 1

Input: A parallel flow graph as defined in Section 2, a program point $v \in N$ and a program variable $y \in X$.

Output: “yes” if y is a copy constant at v ; “no” otherwise.

Method:

- 1) Compute—by standard fixpoint iteration—the least solution over domain (AD, \sqsubseteq) of the constraint system for bridging runs to program point v ; this gives us a value $\text{B}_v^\#[u]$ for each program point u ; as a by-product this computation determines $\text{R}^\#[v]$.
 - 2) Set $I[w] := \{x \mid (1, \langle(x, y)\rangle, 1) \in \text{B}_v^\#[w].2\}$ for each program point $w \in N$.
 - 3) Set $flag := \text{false}$ and $val := \text{unset}$.
 - 4) If $y \in \text{R}^\#[v].1$ or if there is $x \in X$ with $(1, \langle(x, y)\rangle, 1) \in \text{R}^\#[v].2$ then $flag := \text{true}$.
 - 5) For all base edges $e = (u, w)$ annotated by an assignment statement $x := e$ with $x \in I[w]$:
 - 5.1) If e is a composite expression then $flag := \text{true}$;
 - 5.2) If e is a constant expression then
if $val = \text{unset}$ then $val := e$ else if $val \neq e$ then $flag := \text{true}$.
 - 6) If $flag$ then output “no” else output “yes”.
-

Fig. 13. An algorithm that detects copy constants in parallel programs.

Intuitively, $I[w]$ is the set of variables that can influence the value of y at v when some computation is at w . Clearly, in $I[w]$ dependences of bridging runs from w to v are considered. By solving the constraint system for bridging runs from Section 2 over the domain (AD, \sqsubseteq) (Step 1), we can compute the dependence traces of bridging runs; they are given by the second component of the value $\text{B}_v^\#[w]$ that is computed. From the dependence traces we can read off the dependences by Proposition 12 and hence determine $I[w]$ (Step 2). The fixpoint computation in Step 1 determines as a by-product the abstraction $\text{R}^\#[v]$ of the runs reaching v because the constraint system for bridging runs embodies the one for reaching runs.

The rest of the algorithm is based on the following observation: variable y is *not* a copy constant at v if and only if one of the following is true:

- a) there is a variable x the initial value of which can influence y at v ;
- b) there is a base edge $e = (u, w)$ annotated by an assignment $x := e$ with a composite expression e on the right hand side such that x 's value at w can influence y 's value at v ;
- c) there are two distinct base edges $e = (u, w)$ and $e' = (u', w')$ each of them annotated by a constant assignment $x := c$ and $x' := c'$, respectively, such that both x at w and x' at w' can influence y at v and $c \neq c'$.

In Step 3-6 we check whether one of these conditions is true. We use a Boolean variable *flag* that is initialized to false and is set to true once we encounter a reason for *y* not being a copy constant at *v*. Step 4 tests whether condition a) is true: it sets *flag* if the initial value of *y* can flow to *v* ($y \in \mathbb{R}^\#[v].1$) or if the initial value of some variable *x* can influence *y* at *v* via a chain of assignments ($(1, \langle(x, y)\rangle, 1) \in \mathbb{R}^\#[v].2$). Step 5 is concerned with conditions b) and c). Each base edge is examined in turn. Step 5.1 tests whether b) holds. In order to check c), we memorize in a variable *val* the value of the constant assignment that can influence *y* at *w* encountered first. In order to check c) we simply compare the value of constant assignments encountered later with the value memorized in *val*. Variable *val* is initialized with a special value *unset* that indicates that we have not seen a constant assignment so far. Finally, Step 6 outputs the answer.

Of course we could stop the algorithm immediately, once the flag is set to true. Moreover, we can output the value stored in *val* as additional information, if we have identified *y* as a copy constant at *v*. It is the value guaranteed for *y* at *v*. It may happen that *val* has still the value *unset*; this indicates that *v* is an unreachable program point.

We conclude:

Theorem 56 *Algorithm 1 solves the interprocedural copy constant detection problem in parallel flow graphs relative to non-atomic interpretation of base statements.*

5.2 Faint Code Elimination

A variable *x* is *live* at a program point *p* if there is a run from *p* to the end of the program on which *x* is used before it is overwritten. By referring to [9], Horwitz et. al. [32] define a variable *x* as *truly live* at a program point *p* if there is a run from *p* to the end of the program on which *x* is used in a truly live context before being defined, where a truly live context means: in a predicate, or in a call to a library routine, or in an expression whose value is assigned to a truly live variable. True liveness can be seen as a refinement of the ordinary liveness property. We call a use of a variable *x* in a predicate or call to a library routine a *relevant use* of *x*.

Assignments to variables that are not truly live at the program point just after the assignment are called *faint*. Intuitively, faint assignments can not influence any predicate in the program or call of a library routine. Thus, they cannot influence the observable behavior of the program (except of producing run-time errors) and may safely be eliminated from the program. This is called *faint code elimination*.

Algorithm 2

Input: A parallel flow graph as defined in Section 2; a mapping $R : N \rightarrow 2^X$ that associates each program point u with the set of variables relevant at u .

Output: An updated edge annotation A_{new} of the parallel flow graph in which faint code is eliminated.

Method:

- 1) Initialize the new annotation of flow graph edges: $A_{\text{new}} := A$.
 - 2) For each base edge $e \in \text{Base}$: $A_{\text{new}}[e] := \text{skip}$.
 - 3) For each $v \in N$ with $R(v) \neq \emptyset$:
 - 3.1) Compute—by standard fixpoint iteration—the least solution over domain (AD, \sqsubseteq) of the constraint system for bridging runs to program point v ; this gives us a value $B_v^\#[u]$ for each program point u .
 - 3.2) Set $I[w] := \{x \mid \exists y \in R(v) : (1, \langle(x, y)\rangle, 1) \in B_v^\#[u].2\}$ for each program point $w \in N$.
 - 3.3) For each base edge $e = (-, w) \in \text{Base}$ with $A[e] = (x := t)$:
if $x \in I[w]$ then $A_{\text{new}}[e] := (x := t)$.
 - 4) Output the new edge annotation A_{new} .
-

Fig. 14. An algorithm that eliminates faint code in parallel programs.




Fig. 15. A faint assignment that is not dead.

Faint code elimination is a stronger form of the classic transformation of dead-code elimination [7]. Indeed, any assignment that is dead is also faint but not vice versa. The paradigmatic example is shown in Fig. 15. The value computed by $x := x + 1$ in the loop is immediately overwritten after the loop and thus never used in a relevant context. Hence $x := x + 1$ is faint. However, it is not dead because x is potentially (non-relevantly) used by the same statement in the next iteration of the loop. Thus, faint code elimination in general can eliminate more code from a program.

Faint code elimination is based on information about the relevant uses of variables. Typically, this information is derived from the library calls and the conditions in the program. As our view of a program, a parallel flow graph, is an abstraction of the actual program in which library calls as well as conditions are invisible, we assume that we are given this information explicitly in the form of a mapping $R : N \rightarrow 2^X$; for each program point $u \in N$, $R(u)$ is the set of variables directly relevant at u .

Example 57 *In a given program we might find a printf statement, e.g.,*

```
printf("x+y=%d", x+y);
```

In the abstract flow graph view of the program this statement gives rise to a skip edge $e = (u, v)$. Then both x and y are relevant at u , hence $R(u) = \{x, y\}$.

Similarly, we might find a branching statement, e.g.,

```
if (z > x*y) then {...} else {...}
```

In the abstract flow graph view of the program this if-statement gives rise to two skip-edges (u, v) and (u, w) ; u is the start node for the flow graph for the whole if-statement; at v the flow graph for the then part and at w the flow graph for the else part is found. In this case, we have $R(u) = \{x, y, z\}$. \square

Algorithm 2 in Fig. 14 reads a parallel flow graph and a mapping $R : N \rightarrow 2^X$ as described above. Based on this information it calculates an updated version of the edge annotation mapping of the given flow graph in which faint code is eliminated, i.e., faint instances of base statements are replaced by skip.

First the new edge annotation mapping is initialized by the original edge annotation (Step 1) and all annotations of base edges are removed, i.e. replaced by skip (Step 2). The rest of the algorithm restores the original edge annotation for the non-faint base edges. The algorithm is based on the simple idea that an instance of a base statement is not faint if and only if it can influence a relevant value.

We explore all program points v at which at least one variable is relevant and restore the base edges that perform a computation that can influence a variable y that is relevant at v (Step 3). For this purpose we calculate in Steps 3.1 and 3.2 for all program points w the set

$$I[w] = \{x \mid e_{Main} \Longrightarrow c_w \xrightarrow{r} c_v, At_w(c_w), At_v(c_v), \exists y \in R(v) : \hat{r} \text{ exhibits } (x, y)\}.$$

Intuitively, $I[w]$ contains the variables that can influence the value of a relevant variable y at v when some computation is at w . The computation is analogous to the one of the similar set $I[w]$ in Algorithm 1; therefore, we omit a detailed explanation. Step 3.3 restores the annotation of those base edges that assign to a variable that can influence a relevant variable at v from the target node of the base edge. Finally, Step 4 outputs the computed new edge annotation mapping.

We conclude:

Theorem 58 *Algorithm 2 solves the interprocedural faint code elimination problem in parallel flow graphs relative to non-atomic interpretation of base statements.*

5.3 Run-Time

In this section we analyze the asymptotic run-time of the algorithms from the previous sections. We do not determine very sharp estimates but show that the algorithms run in time exponential in the number of program variables, $|X|$, and polynomial in the size of the parallel flow graph. The latter is measured by the parameters $|N|$, the number of program points, $|E|$, the number of edges, and $|\text{Proc}|$, the number of procedures.

In both algorithms the bulk of the work is done in the least fixpoint computation(s). Let us, first of all, determine an asymptotic bound for the complexity of such a fixpoint computation. As we are heading only for a rough bound, we can assume that the least fixpoint is computed naïvely by standard fixpoint iteration: starting from an assignment of the bottom value to each variable appearing in the constraint system we iteratively determine a new assignment to the variables by re-evaluating all constraints until stabilization. Of course the asymptotic complexity of this naive fixpoint algorithm is bounded by the product of the maximal number of iterations and the maximal cost of a single step.

In each iteration except of the last one, at least one constraint variable must change its value. As values only increase during fixpoint iteration, each constraint variable can change its value at most $\mathcal{O}(|X|^{2|X|+2})$ times, because this is a bound for the height of **AD** by Lemma 29. Moreover, it is a simple counting exercise to show that the complete constraint system for bridging runs (it comprises the constraint systems for same-level runs, inverse same-level runs, reaching runs, etc.) has $\mathcal{O}(|\text{Proc}| \cdot |N|)$ constraint variables.⁵ Thus, we can have at most $\mathcal{O}(|\text{Proc}| \cdot |N| \cdot |X|^{2|X|+2})$ iterations. This clearly is $\mathcal{O}(|\text{Proc}| \cdot |N| \cdot 2^{p_0(|X|)})$ for some polynomial $p_0(x)$ in x .

Let us now bound the costs of a single iteration. In each iteration we must reevaluate all constraints. It is again a simple counting exercise to show that the complete constraint system for bridging runs has $\mathcal{O}(|N| \cdot |E|)$ constraints.⁶ From Lemma 54 we know that all operations can be computed in time $\mathcal{O}(2^{p_1(|X|)})$ for some polynomial $p_1(x)$. As the number of operations in each single constraint is bounded, the cost of a single iteration is thus $\mathcal{O}(|N| \cdot |E| \cdot 2^{p_1(|X|)})$.

⁵ This asymptotic bound holds in the special case where ASS1 and ASS2 are true as well as in the general case.

⁶ Again this asymptotic bound holds for both the special and the general case.

Summarizing:

Lemma 59 *The constraint system for bridging runs can be evaluated over domain (AD, \sqsubseteq) in time $\mathcal{O}(|\mathbf{Proc}| \cdot |N|^2 \cdot |E| \cdot 2^{p(|X|)})$ for some polynomial $p(x)$. \square*

Let us now turn attention to the algorithms. Clearly, in the copy constant detection algorithm, Algorithm 1, the bulk of the work is done in Step 1 such that the time taken for Step 1 majorizes the time taken for the other steps. Hence this algorithm runs in time $\mathcal{O}(|\mathbf{Proc}| \cdot |N|^2 \cdot |E| \cdot 2^{p(|X|)})$ by Lemma 59.

In the faint code elimination algorithm, Algorithm 2, the work performed in Step 3.1 majorizes the work done in the other steps. Step 3.1 is executed at most $|N|$ times. Consequently, Algorithm 2 runs in time $\mathcal{O}(|\mathbf{Proc}| \cdot |N|^3 \cdot |E| \cdot 2^{p(|X|)})$.

Clearly, only those program variables are of interest in the algorithms that appear in the parallel flow graph. We can thus assume without loss of generality, that all program variables in X appear in the parallel flow graph. As the latter constitutes part of the input to all algorithm, the input size cannot be smaller than the size of X . Obviously, the same holds for \mathbf{Proc} , N , and E such that the size of the input clearly bounds all the parameters appearing in above run-time estimations. Hence all algorithms run in time exponential in the size of the input.

Theorem 60 *Algorithms 1 and 2 run in exponential time. More precisely, Algorithm 1 runs in time $\mathcal{O}(|\mathbf{Proc}| \cdot |N|^2 \cdot |E| \cdot 2^{p(|X|)})$ and Algorithm 2 in time $\mathcal{O}(|\mathbf{Proc}| \cdot |N|^3 \cdot |E| \cdot 2^{p(|X|)})$. \square*

Corollary 61 *If base statements are interpreted non-atomically, the following two problems can be solved interprocedurally in parallel flow graphs in exponential time: (1) copy constant detection and (2) faint code elimination. \square*

These results raise the question whether there are also *efficient* algorithms for these problems. Sadly, the answer to this question is ‘no’, unless $P=NP$, as we show in the next section.

6 Intractability

We exhibit a co-NP-hardness proof by means of a reduction from the well-known SAT-problem [33,16] that applies to both flow analysis problems. This reduction was first presented in [13] where atomic execution of base statement has been assumed, but it remains valid if this assumption is abandoned. Unlike the reductions in [12] it only relies on propagation along copying assignments

but not on re-initialization. For ease of presentation we represent parallel programs in this section by syntactic programs rather than flow graphs.

6.1 The SAT-Reduction

An instance of SAT is a conjunction $c_1 \wedge \dots \wedge c_k$ of *clauses* c_1, \dots, c_k . Each clause is a disjunction of *literals*; a literal l is either a variable x or a negated variable \bar{x} , where x ranges over some set of variables X . We write $\bar{X} = \{\bar{x}_1, \dots, \bar{x}_n\}$ for the set of negated variables. It is straightforward to define when a *truth assignment* $T : X \rightarrow \mathbb{B}$, where $\mathbb{B} = \{\text{tt}, \text{ff}\}$ is the set of truth values, satisfies $c_1 \wedge \dots \wedge c_k$. The SAT problem asks us to decide for each instance $c_1 \wedge \dots \wedge c_k$ whether there is a satisfying truth assignment or not.

From a given SAT instance $c_1 \wedge \dots \wedge c_k$ with k clauses over n variables $X = \{x_1, \dots, x_n\}$ we construct a loop-free parallel program. In the program we use $k + 1$ variables z_0, z_1, \dots, z_k . Intuitively, validity of clause c_i is related to propagation from z_{i-1} to z_i . For each literal $l \in X \cup \bar{X}$ we define a statement π_l that consists of a sequential composition of assignments of the form $z_i := z_{i-1}$ in increasing order of i . The assignment $z_i := z_{i-1}$ is in π_l if and only if the literal l makes clause i true. Formally, $\pi_l = \pi_l^k$, where

$$\pi_l^0 \stackrel{\text{def}}{=} \mathbf{skip}$$

$$\pi_l^i \stackrel{\text{def}}{=} \begin{cases} \pi_l^{i-1}; z_i := z_{i-1}, & \text{if clause } c_i \text{ contains } l \\ \pi_l^{i-1}, & \text{if clause } c_i \text{ does not contain } l \end{cases}$$

for $i = 1, \dots, k$. Now, consider the following program π , where \parallel denotes non-deterministic choice:

```

procedure Main;
   $z_0 := 1; z_1 := 0; \dots; z_k := 0;$ 
   $[(\pi_{x_1} \parallel \pi_{\bar{x}_1}) \parallel \dots \parallel (\pi_{x_n} \parallel \pi_{\bar{x}_n})];$ 
   $(z_k := 0 \parallel \mathbf{skip}); \mathbf{write}(z_k)$ 
end

```

Clearly, π can be constructed from the given SAT instance $c_1 \wedge \dots \wedge c_k$ in polynomial time or logarithmic space.

It is not hard to see that the value 1 from the initialization of z_0 can be propagated to the final write statement if and only if the given SAT instance is satisfiable:

“**If**”: On the one hand, we can construct from a satisfying truth assignment $T : X \rightarrow \mathbb{B}$ a run that propagates z_0 's initialization to the write-statement.

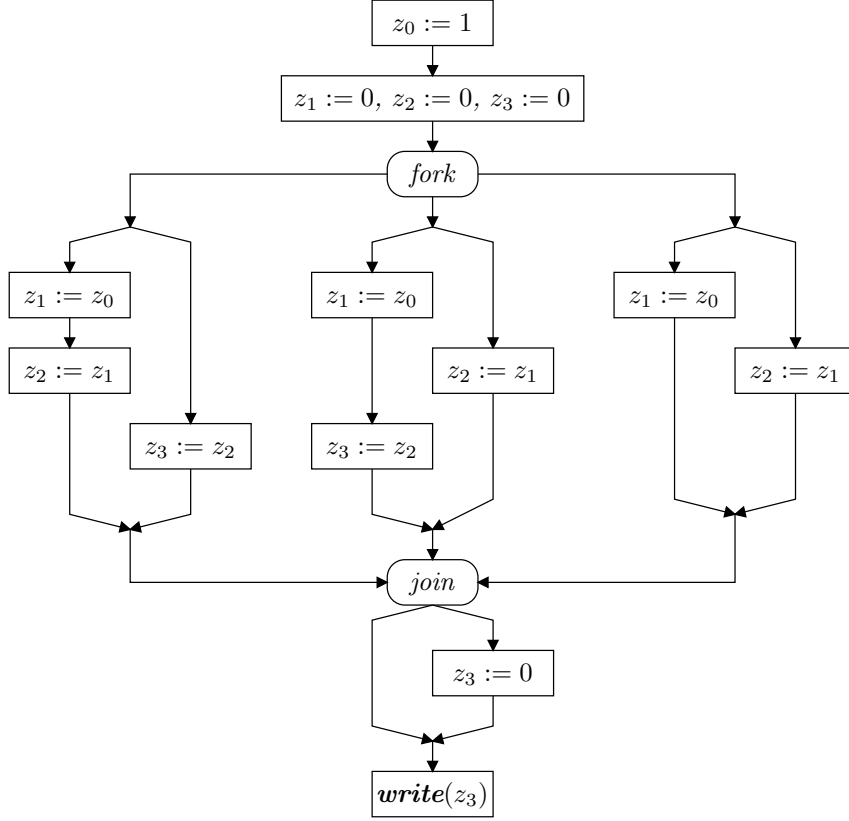


Fig. 16. The flow graph for $(x_1 \vee x_2 \vee x_3) \wedge (x_1 \vee \overline{x_2} \vee \overline{x_3}) \wedge (\overline{x_1} \vee x_2)$.

In each parallel component $\pi_{x_i} \parallel \pi_{\overline{x_i}}$ we choose the left branch π_{x_i} if $T(x_i) = \text{tt}$ and the right branch $\pi_{\overline{x_i}}$ otherwise. As T is a satisfying truth assignment, there will be, for any $i \in \{1, \dots, k\}$, at least one assignment $z_i := z_{i-1}$ in one of the chosen branches. We interleave the branches now in such a way that the assignment(s) to z_1 are executed first, followed by the assignment(s) to z_2 etc. This results in a propagating run.

“Only if”: On the other hand, a run can propagate the initialization value of z_0 to the write-statement only via copying it from z_0 to z_1 , from z_1 to z_2 etc., because all assignments have the form $z_i := z_{i-1}$. Such a run must contain at least one assignment $z_i := z_{i-1}$ for all $i = 1, \dots, k$. From the way in which the non-deterministic choices are resolved in such a run we can easily construct a satisfying truth assignment.

The arguments for both directions hold independently from the atomicity assumption for assignment statements.

Example 62 Fig. 16 shows an example clause and program for illustration. Assignments to different variables are shown on different levels. Intuitively a satisfying truth assignment corresponds to a way of resolving the non-deterministic choices in the three threads such that at each level at least one assignment is present in one of the chosen branches. This is the case if and only if the value

1 from z_0 's initialization may propagate to the write instruction.

It is not hard to infer from this propagation property that the given SAT instance is satisfiable if and only if any of the following two conditions holds:

- (1) $z_0 := 1$ is not a faint assignment.
- (2) z_k is not a copy constant at the write statement.

The second point deserves additional explanation. Observe first that z_k can hold only 0 or 1 at the write-statement because all variables are initialized by 0 or 1 and the other assignments only copy these values. Clearly, due to the non-deterministic choice just before the write-statement, z_k may hold 0 finally. Thus, z_k is a constant at the write-statement if and only if it cannot hold 1 there. The latter obviously holds if and only if the initialization value of z_0 cannot be propagated.

The program constructed in the above reduction is loop-free and does not employ procedures. Therefore, the reduction already applies to the intraprocedural problems for loop-free programs. It is easy to see that the problems can also be solved in non-deterministic polynomial time for loop-free programs: a non-deterministic algorithms may guess two runs that witnesses non-constancy or a single run that witnesses non-faintness, respectively. Each of these runs can visit any program point at most once because the program is loop-free. Hence it can be guessed even in time linear in the program size.

These considerations prove:

Theorem 63 *Independently of the atomicity assumption for base statements, detecting copy constants and detecting faint code in loop-free parallel programs are co-NP-complete problems.*

Corollary 64 *Independently of the atomicity assumption, detecting copy constants, and detecting faint code are co-NP-hard problems in arbitrary parallel programs.*

7 Conclusion

Statements of parallel programs are broken into instructions of the underlying hardware architecture prior to execution. Hence it is unrealistic to assume that statements execute as atomic steps. In this paper, we have shown that relative to non-atomic execution precise interprocedural dependence analysis of parallel programs is possible in sharp contrast to the situation when base statement are assumed to execute atomically. Specifically, we proposed an effective abstract domain of *antichains of dependence traces* which enables a precise

abstract interpretation of constraint systems characterizing (non-atomic) run sets of interest in parallel programs. From the values of this abstract domain the exhibited dependences can be read off.

The dependence traces domain provides us with a means to perform precise interprocedural dependence analysis in parallel programs. We have shown how it can be used for interprocedural detection of copy constants and elimination of faint code. Our algorithms solve the problems completely relative to the non-atomic semantics of parallel programs. Of course, the algorithms are sound also under the stronger execution assumption that base statements execute atomically. However, relative to atomic execution they are incomplete, which is indispensable in view of known undecidability results. The algorithms have exponential worst-case run-time. Indeed, we show that detection of copy constants and faint code elimination remain intractable problems even when the atomic execution idealization is abandoned. This holds already for parallel programs without loops or procedures.

We believe that refinements of the technique underlying dependence traces can lead to practically interesting algorithms that are much more precise than existing algorithms in which interference it treated rather pessimistically. While the run-time of the algorithms is exponential in the number of program variables, it is *polynomial in the program size*. Hence, they are polynomial-time algorithms if the number of program variables is bounded. An interesting direction for future research is to extend the algorithms to treat local variables. Such algorithms should use the expensive dependence traces technique only for tracing propagation via global variables and combine this with a cheap sequential technique for propagation via (thread- or procedure-) local variables. This seems promising because most variables are local in practice.

Acknowledgements

I thank Helmut Seidl for many discussions that stimulated the research reported here and helped to clarify subtle points. He invited me to work for half a year in his research group at the University of Trier which allowed me to elaborate these ideas. Moreover, I am indebted to Jens Knoop, Oliver R uthing, and Bernhard Steffen who shared part of their rich knowledge of data-flow analysis with me.

I thank the anonymous referees of TCS for their very helpful feedback that helped to improve upon the original submission.

References

- [1] M. Rinard, Analysis of multithreaded programs, in: P. Cousot (Ed.), *Static Analysis of Systems (SAS 2001)*, Vol. 2126 of *Lecture Notes in Computer Science*, Springer-Verlag, 2001, pp. 1–19.
- [2] J. Knoop, B. Steffen, J. Vollmer, Parallelism for free: Efficient and optimal bitvector analyses for parallel programs, *ACM Transactions on Programming Languages and Systems* 18 (3) (1996) 268–299.
- [3] J. Knoop, Parallel constant propagation, in: D. Pritchard, J. Reeve (Eds.), *4th European Conference on Parallel Processing (EURO-PAR '98)*, Vol. 1470 of *Lecture Notes in Computer Science*, Springer-Verlag, 1998, pp. 445–455.
- [4] H. Seidl, B. Steffen, Constraint-based inter-procedural analysis of parallel programs, *Nordic Journal of Computing* 7 (4) (2001) 371 – 400.
- [5] J. Esparza, J. Knoop, An automata-theoretic approach to interprocedural data-flow analysis, in: W. Thomas (Ed.), *Foundations of Software Science and Computation Structure (FoSSaCS'99)*, Vol. 1578 of *Lecture Notes in Computer Science*, Springer-Verlag, 1999, pp. 14–30.
- [6] J. Esparza, A. Podelski, Efficient algorithms for pre^* and $post^*$ on interprocedural parallel flow graphs, in: *27th ACM International Conference on Principles of Programming Languages (POPL)*, 2000, pp. 1–11.
- [7] S. S. Muchnick, *Advanced Compiler Design and Implementation*, Morgan Kaufmann Publishers, San Francisco, California, 1997.
- [8] C. Fischer, R. LeBlanc, *Crafting a Compiler*, Benjamin/Cummings Publishing Co., Inc., Menlo Park, CA, 1988.
- [9] R. Giegerich, U. Möncke, R. Wilhelm, Invariance of approximative semantics with respect to program transformations, in: *GI 11. Jahrestagung*, Vol. 50 of *Informatik Fachberichte*, Springer-Verlag, 1981, pp. 1–10.
- [10] M. Weiser, Program slicing, *IEEE Transactions on Software Engineering* 10 (4) (1984) 352–357.
- [11] F. Tip, A survey of program slicing techniques, *Journal of Programming Languages* 3 (3) (1995) 121–189.
- [12] M. Müller-Olm, H. Seidl, On optimal slicing of parallel programs, in: *33th Annual ACM Symposium on Theory of Computing (STOC)*, ACM SIGACT, ACM Press, Hersonissos, Crete, Greece, 2001, pp. 647–656.
- [13] M. Müller-Olm, The complexity of copy constant detection in parallel programs, in: A. Ferreira, H. Reichel (Eds.), *18th Annual Symposium on Theoretical Aspects of Computer Science (STACS 2001)*, Vol. 2010 of *Lecture Notes in Computer Science*, Springer, 2001, pp. 490–501.

- [14] A. Bouajjani, P. Habermehl, Constrained properties, semilinear systems, and Petri nets, in: U. Montantari, V. Sassone (Eds.), 7th International Conference on Concurrency Theory (CONCUR '96), Vol. 1119 of Lecture Notes in Computer Science, Springer-Verlag, 1996, pp. 481–497.
- [15] G. Ramalingam, Context-sensitive synchronization-sensitive analysis is undecidable, *ACM Transactions on Programming Languages and Systems* 22 (2) (2000) 416–430.
- [16] C. H. Papadimitriou, *Computational Complexity*, Addison-Wesley, 1994.
- [17] D. S. Johnson, A catalog of complexity classes, in: J. van Leeuwen (Ed.), *Handbook of Theoretical Computer Science. Volume A: Algorithms and Complexity*, Elsevier Science Publishers B.V., 1990, pp. 67–161.
- [18] F. Nielson, H. R. Nielson, C. Hankin, *Principles of Program Analysis*, Wiley Professional Computing, Springer-Verlag, 1999.
- [19] A. V. Aho, R. Sethi, J. D. Ullman, *Compilers: Principles, Techniques, and Tools*, Addison-Wesley, 1986.
- [20] M. S. Hecht, *Flow Analysis of Computer Programs*, Elsevier North-Holland, 1977.
- [21] P. Cousot, R. Cousot, Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints, in: 4th ACM Symposium on Principles of Programming Languages (POPL), Los Angeles, California, 1977, pp. 238–252.
- [22] P. Cousot, R. Cousot, Abstract interpretation frameworks, *J. Logic Computat.* 4 (2) (1992) 511–547.
- [23] D. Grunwald, H. Srinivasan, Data flow equations for explicitly parallel programs, *SIGPLAN Notices* 28 (7).
- [24] J. C. M. Baeten, W. P. Weijland, *Process Algebra*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, 1990.
- [25] D. Lugiez, P. Schnoebelen, The regular viewpoint on PA-processes, *Theoretical Computer Science* 274 (1-2) (2002) 89–115.
- [26] G. D. Plotkin, A structured approach to operational semantics, Tech. Rep. DAIMI FN-19, Aarhus University, Comput. Sci. Dept. (1981).
- [27] inmos limited, *Transputer Instruction Set – A Compiler Writer’s Guide*, 1st Edition, Prentice Hall International, 1988.
- [28] T. H. Cormen, C. E. Leiserson, R. L. Rivest, *Introduction to Algorithms*, The MIT Press, 1990.
- [29] K. R. Apt, G. D. Plotkin, Countable nondeterminism and random assignment, *Journal of the ACM* 33 (4) (1986) 724–767.

- [30] Mathematics of Program Construction Group, Fixed-point calculus, *Information Processing Letters* 53 (3) (1995) 131–136.
- [31] M. Sagiv, T. Reps, S. Horwitz, Precise interprocedural dataflow analysis with applications to constant propagation, *Theoretical Computer Science* 167 (1–2) (1996) 131–170.
- [32] S. Horwitz, T. Reps, M. Sagiv, Demand interprocedural dataflow analysis, Tech. Rep. TR-1283, Computer Sciences Department, University of Wisconsin, Madison, WI (August 1995).
- [33] S. A. Cook, The complexity of theorem-proving procedures, in: Proc. 3rd Ann. ACM Symp. on Theory of Computing (STOC), 1971, pp. 151–158.