Nordic Journal of Computing

Constraint-Based Inter-Procedural Analysis of Parallel Programs

Helmut Seidl
FB IV – Informatik,
Universität Trier, D-54286 Trier, Germany,
seidl@uni-trier.de

Bernhard Steffen
Lehrstuhl für Programmiersysteme,
Universität Dortmund, Baroper Straße 301,
D-44221 Dortmund, Germany,
Bernhard.Steffen@cs.uni-dortmund.de

Abstract. We provide a uniform framework for the analysis of programs with procedures and explicit, unbounded, fork/join parallelism covering not only bitvector problems like reaching definitions or live variables but also non-bitvector problems like strong copy constant propagation. Due to their structural similarity to the sequential case, the resulting algorithms are as efficient as their widely accepted sequential counterparts, and they can easily be integrated in existing program analysis environments like e.g. MetaFrame or PAG. We are therefore convinced that our method will soon find its way into industrial-scale computer systems.

Keywords: Inter-procedural program analysis, explicit parallelism, bitvector problems, strong copy constant propagation, coincidence theorems.

1. Introduction

The analysis of parallel programs is known as a notoriously hard problem. Even without procedures and with only bounded parallelism the analysis typically suffers from the so-called *state explosion problem*: in general, already the required control structures grow exponentially with the number of parallel components. Bitvector analyses, dominant in most practical compilers, escape this problem in the context of fork/join-parallelism [11, 8]: a simple pre-process is sufficient to adapt sequential intra-procedural bitvector analyses to directly work on *parallel flow graphs* which concisely and explicitly represent the program's parallelism. Key for this adaptation was to change from a *property* analysis (directly associating program points with properties) to an *effect* analysis[1] associating program points with a *property transformer* resembling the effect of the 'preceding' program fragment. The simplicity of the adaption results from the fact that bitvector analyses

[1] Second-order analysis in the terminology of [11].

Received December 20, 2000.

can conceptually be "sliced" into separate analyses for each individual bit-component each of which only requires the consideration of a three-point transformer domain. This setting allowed one to elegantly solve the two central problems inherent in the analysis of parallel programs:

- Interference between parallel components can be treated by checking for the existence of 'killers' in the Gen/Kill terminology in parallel components, which is trivially a 'linear task'.
- Synchronization at join nodes reduces to a simple operation defined on the individual effects of the involved parallel components.

In order to handle also procedures and unbounded parallelism, Esparza and Knoop observed that the described problem profile also admits an automata theoretic treatment [5]. This observation has been carefully developed by Esparza and Podelski in [6]. We will come back to this approach in the second part of this paper, where we discuss the merits of refined program analysis. From the perspective of classical program analysis, however, the constructions in [5] are rather impractical. They require involved automata constructions which return the desired results in a rather indirect way: Instead of determining the abstract values at program points directly, their reachability analysis first computes descriptions of sets of so-called *process trees* corresponding to a particular property at a program point. Thus, for each bit-component of the analysis, a tree automata construction is required which is linear in the *product* of the size of the program and the size of an automaton describing all reachable process trees. The latter automaton can grow linearly in the size of the program as well – implying that the analysis of each component is at least *quadratic* in the program size.

In this paper we present a much more direct framework for the inter-procedural analysis of fork/join parallel programs. We propose a constraint-based approach which naturally arises from an algebraic reformulation of the intra-procedural method presented in [11, 8]. Our approach closely resembles the classical understanding of bitvector analysis, has a complexity which is *linear* in the program size and admits elegant, algebraic proofs.

Summarizing, we contribute to the state of the art by

(1) Providing a uniform characterization of the captured analysis profile which *simultaneously* addresses all involved program entities, e.g., all program variables at once for live variable analysis or all program expressions at once for availability of expressions. Moreover, this profile goes beyond pure bitvector analyses as it e.g. also captures strong copy constant propagation [8].

(2) Basing our development on a constraint characterization of valid parallel execution paths. The constraint system for the actual analyses simply results from an abstract interpretation [3, 4, 2] of this characterization.

(3) Presenting a framework which supports algebraic reasoning. In particular, the proof for prop. 2(3) – resembling the central Main Lemma of [11] – straightforwardly evolves from our profile characterization.

(4) Guaranteeing essentially the same performance as for purely inter-procedural bitvector analyses by exploiting the results of a generalized *possible interference* analysis [11].

As a consequence, the presented framework is tightly tailored for the intended application area. It directly associates the program points with the required information based on classical constraint solving through (e.g., worklist based) fixpoint iteration. This can be exploited to obtain simple implementations in current program analysis generators like DFA& OPT MetaFrame [9] or PAG [1], which provide all the required fixpoint iteration machinery.

The constraint systems discussed up to now are *control structure-oriented*, which implies that joins in the control flow typically impose loss of information. This is common in the 'classical' program optimization scenario, and it guarantees good control over the program size of the optimized program. Motivated by the task to eliminate *all* partial redundancies [12], which is impossible without significantly altering the control structure, [14] proposed a *property-oriented approach* to intra-procedural program analysis, which totally avoids the loss of information at join nodes on the cost of node splitting. The central problem with property-oriented expansion is its potential of code-size explosion, which, e.g. in the case of partial redundancy elimination, in the worst case, can be exponential. In practice, however, this agressive optimization seems to behave much better.
From this perspective, the automata theoretic approach presented in [5, 6] can be regarded as a step towards property-oriented *inter-procedural* program analysis and optimization for fork/join-parallel programs, as it refines the control-structure oriented analysis by separating scenarios for single program locations on the basis of the considered program property. This refined information in terms of so-called *process trees*, which is of no use for classical structure-oriented optimization, may well be used to steer the unfolding of recursion in order to increase the optimization potential. The second part of this paper elaborates on this idea by presenting a constraint system for *execution trees*. Execution trees can be considered as a refinement of process trees, tailored to capture all of the required information for property-oriented *inter-procedural* program analysis of fork/join-parallel programs. It turns out that our approach is at least as efficient as the automata-theoretic approach of [5, 6] in the sense that were are able to proof the complexity bound invisaged but not yet proved for their constructions.
The paper is organized as follows. After formally introducing explicitly parallel programs with procedures in section 2, we define the notion of parallel execution paths in section 3, and specify our analysis problem in section 4. Section 5 then presents a precise effect analysis for procedures, which is the basis for the precise inter-procedural reachability analysis given in section 6. Section 7 discusses possible extensions of our formal development. Section 8 then considers refinements of the analysis questions under consideration. It introduces the concepts of (partial) execution trees and process trees and

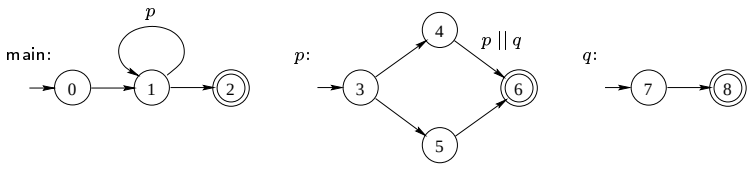

Fig. 1: An Example Control-flow Graph.

provides efficient algorithms to construct regular constraint systems for the result. Finally, section 9 gives our conclusions and perspectives.

2. Programs as Control-flow Graphs

We assume that programs are given as (annotated) control-flow graphs (cfg's for short). An edge in the cfg either is a call of a single procedure, a parallel call to two procedures, or a basic computation step. An example of such a cfg is given in figure 1. There, we only visualized the annotation of call and parallel call edges. Observe that this cfg indeed introduces an unbounded number of instances of procedure q running in parallel.
Formally, a control-flow graph $\mathcal{G}$ for a program with procedures and explicit parallelism consists of a finite set Proc of procedures together with a collection $\mathcal{G}_p, p \in \mathsf{Proc}$, of disjoint *intra-procedural* control-flow graphs. We assume that there is one special procedure main with which program execution starts. The intra-procedural control-flow graph $\mathcal{G}_p$ of a procedure p consists of:

- A set N_p of *program points*;
- A special *entry point* $s \in N_p$ as well as a special *return point* $r \in N_p$;
- A set of edges $E_p \subseteq N_p \times N_p$;
- A subset $C_p \subseteq E_p$ of *call edges* where for $e \in C_p$, $\mathsf{call}\, e = p$ denotes that edge e calls the procedure p; and finally,
- A subset $P_p \subseteq E_p$ of *parallel call edges* where for $e \in P_p$, $\mathsf{call}\, e = p_1 \,||\, p_2$ denotes that edge e calls the procedures p_1 and p_2 in parallel.

Edges which are not contained in C_p or P_p are also called *basic edges*.

Practical Remark: It is just for convenience that we allow only binary parallelism in our programs. Our methods can be easily adapted to work also for more procedures being called in parallel or even parallel do-loops. Also note that we do not consider synchronization between parallel threads by barriers or semaphores. Such constructs *limit* the amount of possible execution paths. By ignoring these, we may get *more* possible execution paths and thus (perhaps less precise but) still *safe* analysis results.

3. Parallel Execution Paths

The semantics of a parallel program is determined w.r.t. the set of parallel execution paths. What we are now going to formalize is an *interleaving semantics* for parallely executable threads. We need the following auxiliary definitions. Let E denote a finite set of edges. Let $w = e_1 \dots e_n$ be a word from E^* and $I = \{i_1 < \dots < i_k\} \subseteq \{1, \dots, n\}$ be a subset of positions in w. Then the *restriction* of w to I is given by $w|_I = e_{i_1} \dots e_{i_k}$.
The *interleaving* of subsets $M_1, M_2 \subseteq E^*$ is defined by

$$M_1 \otimes M_2 = \{w \in E^* \mid \exists I_1 + I_2 = \{1, \dots, |w|\} : w|_{I_1} \in M_1 \text{ and } w|_{I_2} \in M_2\}$$

Here, "+" denotes disjoint union of sets. Thus, $M_1 \otimes M_2$ consists of all possible interleavings of sequences from M_1 and M_2. Furthermore for $M \subseteq E^*$, let $\mathrm{pre}(M)$ denote the set of all prefixes of words in M, i.e.,

$$\mathrm{pre}(M) = \{u \in E^* \mid \exists v \in E^* : uv \in M\}$$

We consider the following sets of possible execution paths:

- For $p \in \mathsf{Proc}$, the set $\Pi(p)$ of all execution paths for p;
- For program point v of procedure p, the set $\Pi(v)$ of all paths starting at the entry point of p and reaching v on the *same level*, i.e., through a path which already has returned from every subsequently entered procedure.
- For every procedure p, the set $\Pi_r(p)$ of all paths starting at a call of main and *reaching* some call of p;
- For every program point v, the set $\Pi_r(v)$ of all paths starting at a call of main and *reaching* program point v.

These sets are given through the least solutions of the following constraint systems (whose variables for simplicity are denoted by $\Pi(p), \Pi(v), \Pi_r(p), \Pi_r(v)$ as well). Let us start with the defining constraint system for the sets of same-level execution paths.

$$\begin{array}{lcll}
\Pi(p) & \supseteq & \Pi(r) & r \text{ return point of } p \quad (1)\\
\Pi(s) & \supseteq & \{\epsilon\} & s \text{ entry point of a procedure} \quad (2)\\
\Pi(v) & \supseteq & \Pi(u) \cdot \{e\} & e = (u, v) \text{ basic edge} \quad (3)\\
\Pi(v) & \supseteq & \Pi(u) \cdot \Pi(p) & e = (u, v) \text{ calls } p \quad (4)\\
\Pi(v) & \supseteq & \Pi(u) \cdot (\Pi(p_1) \otimes \Pi(p_2)) & e = (u, v) \text{ calls } p_1 \,||\, p_2 \quad (5)
\end{array}$$

Here, the binary operator "$\cdot$" denotes concatenation of sequences extended to sets of sequences.
Lines (1) through (4) are the standard lines to determine the sets of all same-level execution paths as known from inter-procedural analysis of sequential programs. Line (1) says that the set of execution paths of procedure p is the set of same-level paths reaching the return point of p. Line (2) says that at least ϵ is a same-level execution path that reaches the entry point

of a procedure. Line (3) says that for every basic edge $e = (u, v)$, the set of same-level execution paths reaching the program point v subsumes all same-level execution paths to u extended by e. Line (4) says that for every edge $e = (u, v)$ calling a procedure p, the set of same-level execution paths reaching the program point v subsumes all same-level execution paths reaching u extended by any execution path through the procedure p. Line (5) for a parallel call of $p_1 \,||\, p_2$ has the same form as line (4). But now the same-level execution paths to the program point before the call are extended by all interleavings of execution paths for p_1 and p_2.
In order to specify the sets $\Pi_r(p), \Pi_r(v)$, let us introduce the auxiliary sets $\Pi(v, p)$, v a program point, p a procedure, which give the sets of execution paths reaching v from a call of p. These auxiliary sets are defined as the least solution of the following system of constraints:

$$\begin{array}{lcll} \Pi(v,p) & \supseteq & \Pi(v) & v \text{ program point of procedure } p \quad (1) \\ \Pi(v,p) & \supseteq & \Pi(u) \cdot \Pi(v,q) & e = (u,_) \text{ calls } q \text{ in } p \quad (2) \\ \Pi(v,p) & \supseteq & \Pi(u) \cdot (\Pi(v,p_i) \otimes M) & e = (u,_) \text{ calls } p_1 \,||\, p_2 \text{ in } p \quad (3) \end{array}$$

where M in line (3) is given by $M = \mathrm{pre}(\Pi(p_{3-i}))$. The intuition behind this definition is as follows. Line (1) says that whenever v is a program point of procedure p, then the set of execution paths from p to v subsumes all same-level execution paths from p to v. Line (2) says that whenever at some edge $e = (u, _)$ in the body of procedure p, some procedure q is called, then the set of execution paths from p to v subsumes all computation paths consisting of a same-level execution path from p to the program point u followed by an execution path from q to v. Finally, line (3) considers an edge $e = (u, _)$ in the body of p which is a parallel call of p_1 and p_2. Then we have to append to the same-level execution paths to u all interleavings of execution paths from p_i to v with prefixes of same-level execution paths for the parallel procedure.
Given the $\Pi(v, q)$, we define the values $\Pi_r(v), \Pi_r(p)$ as the least solution of:

$$\begin{array}{lcll} \Pi_r(v) & \supseteq & \Pi(v, \mathsf{main}) & v \text{ a program point} \\ \Pi_r(p) & \supseteq & \Pi_r(u) & \text{edge } (u,_) \text{ calls } p,\ p \,||\, _ \text{ or } _ \,||\, p \end{array}$$

For now, let us assume that all the sets of execution paths $\Pi(v), \Pi_r(v), \Pi(p), \Pi_r(p)$ are non-empty. In section 7 we will explain how this assumption can be removed.

4. Semantics

Let $\mathbb{D}$ denote a complete lattice and $\mathbb{F} \subseteq \mathbb{D} \to \mathbb{D}$ a subset of monotonic functions from $\mathbb{D}$ to $\mathbb{D}$ which contains $\lambda x.\bot$ (the constant $\bot$-function) and $I = \lambda x.x$ (the identity) and is closed under composition "$\circ$" and least upper bounds. While $\mathbb{D}$ is meant to specify the set of abstract properties, $\mathbb{F}$ describes all possible ways how properties may be transformed when passing from one program point to the other. In this paper we make the following additional assumption:

- $\mathbb{D}$ is *distributive*, i.e., $a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c)$ holds for all $a, b, c \in \mathbb{D}$;
- $\mathbb{D}$ has height $h < \infty$, i.e., every ascending chain of elements in $\mathbb{D}$ has length at most $h + 1$;
- set $\mathbb{F}$ consists of all functions of the form $f\, x = (a \sqcap x) \sqcup b$ with $a, b \in \mathbb{D}$.

Since $\mathbb{D}$ is distributive, all functions f in $\mathbb{F}$ are distributive as well, i.e., $f\,(a \sqcup b) = (f\,a) \sqcup (f\,b)$ for all $a, b \in \mathbb{D}$. Let us also mention that neither $\mathbb{D}$ nor $\mathbb{F}$ is demanded to be finite. However, since $\mathbb{F}$ is a projection of $\mathbb{D} \times \mathbb{D}$ and $\mathbb{D}$ has height h, the lattice $\mathbb{F}$ has height at most $2h$.
The most prominent class of problems that satisfy our restrictions are *bitvector problems* like available expressions, reaching definitions, life variables or very busy expressions [7]. In these cases, we may choose $\mathbb{D} = \mathbb{B}^h$ where $\mathbb{B} = \{0 \sqsubset 1\}$.
There are, however, further analysis problems which meet our assumptions without being bitvector problems. This is the case, e.g., for strong copy constant propagation. Strong copy constant propagation tries to determine whether or not some constant has been assigned to a variable which later-on remains unchanged. For this application, we may choose $\mathbb{D} = V \to \mathbb{B}$ where V is the set of program variables and $\mathbb{B}$ is the flat lattice of possible values for program variables. Thus, an abstract value $d \in \mathbb{D}$ represents an assignment of program variables to values. In particular, $\mathbb{D}$ has height $h = 2 \cdot \#V$. Note furthermore that for strong copy constant propagation, all functions $f \in \mathbb{F}$ operate on the abstract value for each program variable separately. Each such component-wise transfer function is of the special form $\lambda x.(a \sqcap x) \sqcup b$ with $a \in \{\bot, \top\}$.
Consider, e.g., an assignment statement $X = 7;$ (X a program variable). Then the component-wise transfer functions f_Y for the corresponding edge in the control-flow graph are given by:

$$\begin{array}{lcll} f_X & = & \lambda x.7 & = & \lambda x.(\bot \sqcap x) \sqcup 7 \\ f_Y & = & \lambda x.x & = & \lambda x.(\top \sqcap x) \qquad\qquad \text{for } Y \neq X \end{array}$$

We conclude that ascending chains of functions for strong copy constant propagation have length at most $3 \cdot \#V$.
Let E denote a set of edges and $[.] : E \to \mathbb{F}$ denote an assignment of functions to edges. Then we extend $[.]$ to sequences $w = e_1 \ldots e_n \in E^*$ and sets $M \subseteq E^*$ in the natural way, i.e., by

$$[w] = [e_n] \circ \ldots \circ [e_1] \qquad\qquad [M] = \bigsqcup \{[w] \mid w \in M\}$$

Thus, especially, $[\emptyset] = \lambda x.\bot$ (the least element in $\mathbb{F}$), and $[\{\epsilon\}] = [\epsilon] = I$. Functions $[w]$ and $[M]$ are also called the *effect* of the sequence w and the set M, respectively.
For the rest of this paper we assume that we are given an assignment

$$[e] = f_e = \lambda x.(a_e \sqcap x) \sqcup b_e \in \mathbb{F}$$

to each basic edge e of our input program. Then program analysis tries to compute (approximations to) the following values:

Effects of Procedures: For each procedure p, $\mathsf{Effect}(p) := [\Pi(p)]$ denotes the effect of the set of all same-level execution paths through p;

Reachability: For a start value $d_0 \in \mathbb{D}$, program point v and procedure p, $\mathsf{Reach}(v) := [\Pi_r(v)]\, d_0$ and $\mathsf{Reach}(p) := [\Pi_r(p)]\, d_0$ denote the least upper bounds on all abstract values reaching v along execution paths from main and the least upper bound on all abstract values reaching calls to p, respectively.

The system of these values is called the *Merge-Over-all-Paths* solution (abbreviated: MOP solution) of the analysis problem. Since the respective sets of execution paths are typically infinite, it is not clear whether this solution can be computed effectively. The standard approach proposed in data-flow analysis and abstract interpretation [3, 4, 2] consists in putting up a set $\mathcal{C}$ of constraints on the values we are interested in. The constraints are chosen in such a way that any solution to $\mathcal{C}$ is guaranteed to represent a safe approximation of the values. Quite frequently, however, the least solution of $\mathcal{C}$ equals the MOP solution [10, 13]. Then we speak of *coincidence* of the solutions, meaning that $\mathcal{C}$ *precisely* characterizes the MOP.
In our present application, we are already given a constraint system whose least solution represents the sets of execution paths which are to be evaluated. By inspecting this constraint system, we would naturally try to obtain constraint systems for effect analysis and reachability just by abstracting the lattice of sets of paths with our lattice $\mathbb{F}$. Thus, the ordering "$\subseteq$" of set inclusion on sets of paths is mapped to the ordering on $\mathbb{F}$; set union and concatenation is mapped to least upper bounds and composition of functions. Indeed, this abstraction mapping [.] has the following properties:

PROPOSITION 1. *Let $M_1, M_2 \subseteq E^*$. Then the following holds:*
(1) $[M_1 \cup M_2] = [M_1] \sqcup [M_2]$;
(2) $[M_1 \cdot M_2] = [M_2] \circ [M_1]$ *if both M_1 and M_2 are non-empty.* □

Proposition 1 suggests a direct translation of the constraint system for the sets of execution paths into a constraint system which we are aiming at. The only two obstacles withstanding a direct translation are an abstract interleaving operator (which for simplicity is denoted by "$\otimes$" as well), and a way how to deal with prefixes. For our abstract lattices, these two problems turn out to have surprisingly simple solutions.
For $f_i = \lambda x.(a_i \sqcap x) \sqcup b_i$, $i = 1, 2$, we define the *interleaving* of f_1 and f_2 by:

$$f_1 \otimes f_2 = \lambda x.(a_1 \sqcap a_2 \sqcap x) \sqcup b_1 \sqcup b_2$$

We have:

PROPOSITION 2. *Let $f_1, f_2, f \in \mathbb{F}$. Then the following holds:*
(1) $f_1 \otimes f_2 = f_1 \circ f_2 \sqcup f_2 \circ f_1$;
(2) $(f_1 \sqcup f_2) \otimes f = f_1 \otimes f \sqcup f_2 \otimes f$;

(3) $[M_1 \otimes M_2] = [M_1] \otimes [M_2]$ *for non-empty subsets* $M_1, M_2 \subseteq E^*$.

For a proof of Proposition 2 see Appendix A. In particular, statement (1) of Proposition 2 can be interpreted that, instead of considering *all* possible interleavings of parallel components, it suffices to consider all possibilities of their *sequential* compositions.

5. Effect Analysis

Now we have all prerequisites together to present a constraint system for effect analysis. The least solution of the constraint system defines values $[p]$ for the effect of procedures p together with values $[v]$ for the effects of same-level execution paths reaching program point v.

$[p] \sqsupseteq [r]$	r return point of p	(1)
$[s] \sqsupseteq I$	s entry point	(2)
$[v] \sqsupseteq f_e \circ [u]$	$e = (u, v)$ basic edge	(3)
$[v] \sqsupseteq [p] \circ [u]$	$e = (u, v)$ calls p	(4)
$[v] \sqsupseteq ([p_1] \otimes [p_2]) \circ [u]$	$e = (u, v)$ calls $p_1 \,\|\|\, p_2$	(5)

Lines (1) through (4) are the lines to determine the effects of procedures as known from inter-procedural analysis of sequential programs. Line (1) says that the effect of procedure p is the effect of what has been accumulated for the return point of p. Line (2) says that accumulation of effects starts at entry points of procedures with the identity function $I = \lambda x.x$. Line (3) says that the contribution of a basic edge $e = (u, v)$ to the value for v is given by the value for u extended by the application of the function f_e associated with this edge. Line (4) says that the contribution of an edge $e = (u, v)$ calling a procedure p is determined analogously with the only difference that the function f_e in line (3) is now replaced with the effect $[p]$ of the called procedure. Also line (5) for a parallel call has the same form. But now, in order to determine the combined effect of the parallely executed procedures p_1 and p_2, we rely on the interleaving operator "$\otimes$". This constraint system for effect analysis is the direct abstraction of the corresponding constraint system for same-level reaching paths from section 3. Therefore, we obtain (by distributivity of all involved operators):

THEOREM 1. *The least solution of the effect constraint system precisely describes the effect of procedures, i.e.,*

$$\mathsf{Effect}(p) = [p] \qquad \textit{and} \qquad \mathsf{Effect}(v) = [v]$$

for every procedure p and program point v. These values can be computed in time $\mathcal{O}(h \cdot n)$ where h is the height of the complete lattice, and n is the size of the program. □

6. A Constraint System for Reachability

As for effect analysis, we could mimic the least fixpoint definition of the sets of reaching execution paths through a corresponding constraint system over $\mathbb{F}$. Observe, however, that our defining constraint system for reaching execution paths in section 3 has quadratic size. Clearly, we would like to improve on this, and indeed this is possible – even without sacrificing precision.
Consider the set $\mathrm{pre}(M)$ of prefixes of a non-empty set $M \subseteq E^*$. We make the following key observation:

PROPOSITION 3. *Let E_M denote the edges occurring in elements of M where for $e \in E_M$, $[e] = \lambda x.(a_e \sqcap x) \sqcup b_e$. Then*

$$[pre(M)] = \lambda x.x \sqcup B \qquad \textit{where} \qquad B = \bigsqcup\{b_e \mid e \in E_M\} \qquad \square$$

Thus, all the intersections with the a_e have disappeared. What only remains is the least upper bound on the values b_e. A proof of proposition 3 can be found in Appendix B.
Instead of accumulating effects in a topdown fashion as was necessary in the precise definition of reaching execution paths, we prefer a bottom-up accumulation – a strategy which is commonly used in inter-procedural analysis of sequential programs. There, accumulation directly starts at the main program and then successively proceeds to called procedures.
For each program point v, let $B(v)$ denote the least upper bound of all b_e, for all basic edges e possibly executed in parallel with v. This value is also called *possible interference* of v. Furthermore, We use the auxiliary values $\sigma(p)$, p a procedure, to calculate the least upper bound on b_e for all basic edges possibly executed during evaluation of p. These values are determined through the least solution of the following constraint system:

$\sigma(p)$	$\sqsupseteq$	b_e	e basic edge in procedure p
$\sigma(p)$	$\sqsupseteq$	$\sigma(q)$	procedure p calls q or $q \,\|\|\, _$ or $_ \,\|\|\, q$
$B(v)$	$\sqsupseteq$	$B(p)$	v program point in p
$B(p)$	$\sqsupseteq$	$B(u)$	$(u, _)$ calls procedure p
$B(p_1)$	$\sqsupseteq$	$\sigma(p_2) \sqcup B(u)$	$(u, _)$ calls $p_1 \,\|\|\, p_2$
$B(p_2)$	$\sqsupseteq$	$\sigma(p_1) \sqcup B(u)$	$(u, _)$ calls $p_1 \,\|\|\, p_2$

The whole system for computing the values $\sigma(p)$, $B(p)$ and $B(v)$ is of linear size and uses "$\sqcup$" as only operation in right-hand sides. Such kind of problems are also known as "pure merge problems" and can be solved even in linear time.
We will now construct a constraint system as for inter-procedural reachability analysis of sequential programs, but for each program point additionally take its possible interference into account. Thus, we consider the values $[\![v]\!]$, v a program point, $[\![p]\!]$, p a procedure, which are determined as the least

solution of the following constraint system:

$$
\begin{array}{lll}
[\![\mathsf{main}]\!] \sqsupseteq d_0 & & (1)\\
[\![v]\!] \sqsupseteq B(v) & \text{and} & (2)\\
[\![v]\!] \sqsupseteq [v][\![p]\!] & v \text{ program point in procedure } p & (3)\\
[\![p]\!] \sqsupseteq [\![u]\!] & e = (u, _) \text{ calls } p \text{ or } p \,||\, _ \text{ or } _ \,||\, p & (4)
\end{array}
$$

Only line (2) makes the difference to a corresponding constraint system for reachability in sequential programs. The intuition behind the constraint system is as follows. Line (1) says that initially the value reaching main should subsume the initial value d_0. Line (2) says that the value reaching program point v should subsume its possible interference. Line (3) says that when v is a program point of procedure p, then the reaching value should also subsume the intra-procedural effect of v applied to the value reaching p. Line (4) finally says that the value reaching a procedure should subsume the value of every program point where such a call (possibly in parallel to another call) is possible.
This constraint system differs considerably from the constraint system for the sets of reaching execution paths. Nonetheless, we are able to prove:

THEOREM 2. *The above constraint system computes precise reachability information as well, i.e.,*

$$\mathsf{Reach}(p) = [\![p]\!] \qquad \textit{and} \qquad \mathsf{Reach}(v) = [\![v]\!]$$

for all program points v and procedures p. These values can be computed in time $\mathcal{O}(h \cdot n)$ where h is the height of the complete lattice, and n is the size of the program.

For a proof see Appendix C. Theorem 2 implies that programs with procedures and parallelism are not harder to analyze than programs with procedures but without parallelism!

7. Extensions

In this section, we discuss issues which are important for the practical applicability of the presented results. We do not claim that this section contains any new ideas or constructions. Rather we want to emphasize that the constructions known from the inter-procedural analysis of sequential programs can be extended to parallel programs in a straight-forward way.

7.1 Non-Reachable Program Points

So far, we assumed that every program point is reachable by at least one execution path. In order to show that this assumption is not vital, we now abandon this assumption. Instead, let $\mathcal{P}$ and $\mathcal{R}$ denote the sets of possibly terminating procedures and reachable program points, respectively. In order

to compute these sets, we instantiate our generic analysis with $\mathbb{D} = \{0 \sqsubset 1\}$ where for each basic edge e, the function $[e] = f_e$ is given by $f_e = I = \lambda x.x$, and the initial value d_0 equals 1. The only functions from $\mathbb{D} \to \mathbb{D}$ occurring during the analysis are $\lambda x.0$ and $\lambda x.x$. Both functions are *strict*, i.e., preserve the least element. Therefore, we obtain:

PROPOSITION 4. *For every procedure p and program point v, the following holds:*

(1) $[v] = I$ *iff* $\Pi(v) \neq \emptyset$ *and* $[p] = I$ *iff* $\Pi(p) \neq \emptyset$*;*

(2) $[\![v]\!] = 1$ *iff* $\Pi_r(v) \neq \emptyset$ *and* $[\![p]\!] = 1$ *iff* $\Pi_r(p) \neq \emptyset$*.*

In particular, $p \in \mathcal{P}$ *iff* $[p] = I$*, and* $v \in \mathcal{R}$ *iff* $[\![v]\!] = 1$*.* □

We conclude that the sets $\mathcal{P}$ and $\mathcal{R}$ can be computed in linear time.

A non-reachable program point should not influence any other program point. Therefore, we modify the given cfg by removing all edges starting in program points not in $\mathcal{R}$. By this edge removal, the sets of reaching execution paths have not changed. Let us call the resulting cfg *normalized*. Then we obtain:

THEOREM 3. *Assume the cfg is normalized. Then for every program point v and procedure p,*

(1) $\mathsf{Effect}(v) = [v]$ *and* $\mathsf{Effect}(p) = [p]$*;*

(2) $\mathsf{Reach}(v) = [\![v]\!]$ *and* $\mathsf{Reach}(p) = [\![p]\!]$*.* □

We conclude that, after the preprocessing step of normalization, our constraint systems will compute a safe approximation which is precise.

Practical Remark: Normalization of the cfg may remove edges and thus some constraints from the constraint systems of the analysis. Therefore, omitting normalization may result in a less precise, but still safe analysis.

7.2 Backward Analysis

What we discussed so far, is called *forward analysis*. Examples of forward analysis problems are reaching definitions, available expressions or strong copy constant propagation. Other important analyses, however, determine the value at a program point v w.r.t. the *possible future* of v, i.e., the set of reverses of execution paths possibly following a visit of v. Examples are live variables or very busy expressions. Such analyses are called *backward analyses*. In case that every forward reachable program point is also backward reachable, i.e., lies on an execution path from the start point to the return point of `main`, we can reduce backward analysis to forward analysis – simply by normalizing the cfg followed by a reversal of edge orientations and an exchange of entry and return points of procedures.

7.3 Local and Global State

Consider an edge $e = (u, v)$ in the cfg which calls a terminating procedure p (the treatment of a terminating parallel call to two procedures p_1 and p_2 is completely analogous). So far, the complete information at program point u is passed to the entry point of p. Indeed, this is adequate when analyzing *global* properties like availability of expressions which depend on global variables only. It is not (immediately) applicable in presence of *local* variables which are visible to the caller but should be hidden from the callee p, meaning that they should survive the call unchanged [10, 13].
To make things precise, let us assume that $\mathbb{D} = \mathbb{D}_l \times \mathbb{D}_g$ where $\mathbb{D}_l$ and $\mathbb{D}_g$ describe local and global properties, respectively. Let us further assume that the global part of the current state is passed as a parameter to p, and also returned as the result of the call, whereas the local part of the program point before the call is by-passed the call using some transformer $\beta_e : \mathbb{D}_l \to \mathbb{D}_l$. Recall that every $f \in \mathbb{F}$ is of the form $f = \lambda x.(x \sqcap a) \sqcup b$ with $a, b \in \mathbb{D}$. Since $\mathbb{D}$ is a Cartesian product, this implies that $f = f_l \times f_g$ where $f_l : \mathbb{D}_l \to \mathbb{D}_l$ and $f_g : \mathbb{D}_g \to \mathbb{D}_g$ independently operate on the local states and global states, respectively. Obviously, such an independence will not hold in general for inter-procedural analyses. Instead, it is a very specific consequence from the special form of transfer functions under consideration. Due to this independence, however, we can separate the analysis into two phases.
The first phase considers just global values from $\mathbb{D}_g$. No local state need to be preserved during the call, and we use the original call edge.
The second phase then is purely intra-procedural and deals with the lattice $\mathbb{D}_l$. But now, since the call at edge e has no effect onto the local state, we simply change e into a basic edge with $[e] = \beta_e$.

8. Refining Analyses

Up to now, we were concentrating on safe approximations capturing *all* values reaching a certain program point, i.e. a certain position in the argument program's code. The *refined analysis* we are going to dicuss in this section avoids to collapse, and therefore to weaken, the information belonging to a program point by qualifying the computed information with characterizations of the relevant context.
This can be seen as a generalization of the *property-oriented expansion* approach of [14]. There, the loss of information at join nodes was avoided by node splitting: whenever a 'meet' operation would lead to a proper loss of information, the target node is split in order to be able to maintain the full information separately for the different contexts. The power of this approach becomes apparent, when looking at *partial redundancy elimination*: almost automatically, this approach leads to the elimination of *all* partial redundancies.
In the following we are elaborating on this idea in the setting of the inter-

procedural analysis of fork/join parallel programs. In this setting it is no longer possible to characterize the relevant contexts simply by means of separating versions of the considered program points. Rather, the characterizations must additionally keep track of the nesting of procedure calls and the state of the concurrent computations. Technically, the extra information about the program configuration in which an abstract value may occur, will be represented as a *tree*. A convenient way of representing possibly infinite sets of trees is provided by finite tree automata or *regular constraint systems*. Regular constraint systems ensure that there are efficient algorithms for basic problems like determining the set of constraint variables for which the corresponding sets in the least solution

- ... are non-empty;
- ... contain a specific tree;
- ... are contained in some regular set;
- ... have a non-empty intersection with some regular set.

For **refined reachability** as considered in [5, 6], we are interested in determining for all program points u and abstract values $d \in \mathbb{D}$, the sets of all *process trees* possibly reaching u such that *all* corresponding program executions result in values bounded by d. Here, process trees are tree-like generalizations of call stacks describing the set of currently active parallel processes. The generalization of property-oriented expansion requires even more precise context information. For all program points u and abstract values $d \in \mathbb{D}$, we will therefore compute the set of all *program executions* reaching u and resulting in values bounded by d.

8.1 Execution Trees and Process Trees

Program execution paths (as defined in section 3) do not capture the information about the nesting of procedure calls or the state of parallel executions. Thus they are insufficient to provide a basis for our refined analysis. We therefore consider the enriched structure of *execution trees*, which is designed to capture the required information. An example execution tree for the program from fig. 1, is shown in fig. 2. Intuitively, executions are represented in these tree as (left to right depths first) pre-order tree traversals. Basic execution steps are represented as nodes with one predecessor. Binary nodes are then introduced for procedure calls. The first subtree at such a node represents the execution before the call whereas the second subtree gives the execution of the called procedure. In the example tree of fig. 2, execution of procedure main starts at the leftmost leaf with basic edge $(0, 1)$ followed by a call to procedure p at edge $(1, 1)$. Consequently, the right subtree at $(1, 1)$ represents an execution of p. This execution again starts at the leftmost leaf of the corresponding subtree, i.e., now with the basic edge $(3, 4)$ followed by a parallel call to the procedures p and q at edge $(4, 6)$.
Fig. 2 represents a *full* execution of the considered example program. However, the same structure is easily adapted to also capture the notion of

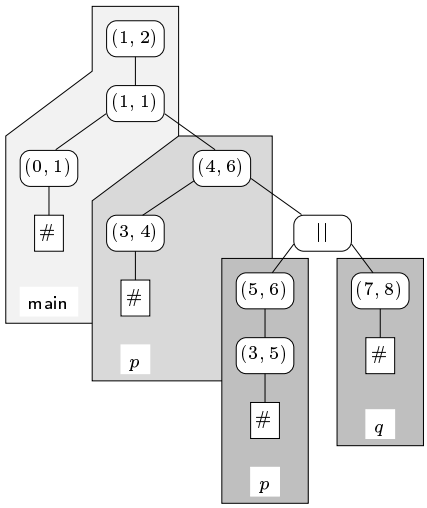

Fig. 2: An Execution Tree for main.

'program point in a certain context'. Accordingly, we also use trees for describing program executions which *reach* a certain program point. An example of such a tree is given in fig. 3.
The idea is to introduce extra binary nodes (here: the barred nodes $(1,1)$ and $(4,6)$) for calls whose evaluation is started but will not be completed before reaching the program point in question. The resulting tree is called *partial execution tree*. The program point reached by a partial execution tree is the program point reach inside the last procedure of the process currently under "focus". Our convention is that the procedures entered by the process under focus are identified by a leftmost leaf $\bullet$. Thus, in fig. 3 the call to the

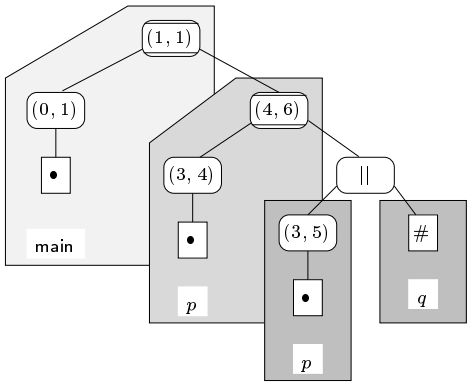

Fig. 3: A Partial Execution Tree Reaching Node 5.

procedure main, has been entered followed by two calls to procedure p where the current program point is given by the endpoint of edge $(3,5)$, i.e., 5. In parallel to this second call to p, a call to procedure q has been spawned. Since the corresponding subtree is the empty tree #, we conclude that (in the represented scenario) so far no edge has been traversed inside q.
The advantage of using such tree-like representations is apparent: since the nesting structure of calls is made explicit, we are able to restore (parts of) the current local state after a completed call.
For a formal treatment, we interpret the set of edges E as a ranked alphabet where basic edges and call edges have ranks 1 and 2, respectively. Furthermore, we need nullary symbols # and $\bullet$ denoting the empty tree and the empty partial tree, respectively, a binary symbol $||$ (written in infix notation) and, for each call edge e a new symbol $\bar{e}$ of rank 2 denoting a call which has been entered but not yet completed. Then, (formal parallel execution) trees t and partial trees r are represented as expression trees as follows:

$$\begin{array}{lll} t & ::= & \# \mid e(t) \mid e(t,t_1) \mid e(t, t_1 \,||\, t_2) \\ r & ::= & \bullet \mid \bar{e}(r,r) \mid \bar{e}(r, r_1 \,||\, t) \mid \bar{e}(r, t \,||\, r_1) \mid e(r) \mid e(r,t) \mid e(r, t_1 \,||\, t_2) \end{array}$$

Each (partial) execution tree x represents a set $\mathcal{I}[x]$ of possible parallel execution sequences:

$$\begin{array}{lllll} \mathcal{I}[\bullet] & = & \mathcal{I}[\#] & = & \{\epsilon\} \\ \mathcal{I}[e(x)] & = & \mathcal{I}[x] \cdot e & & \\ \mathcal{I}[e(x_1,x_2)] & = & \mathcal{I}[\bar{e}(x_1,x_2)] & = & \mathcal{I}[x_1] \cdot \mathcal{I}[x_2] \\ \mathcal{I}[x_1 \,||\, x_2] & = & \mathcal{I}[x_1] \otimes \mathcal{I}[x_2] & & \end{array}$$

The mapping $\mathcal{I}$ is extended to sets of (partial) trees by:

$$\mathcal{I}[X] \;=\; \bigcup \{ \mathcal{I}[x] \mid x \in X \}$$

For procedures p and program points u, we consider the following sets of trees:

- $T[p]$ and $T[u]$ are the sets of parallel execution trees of p and u, respectively;
- $R[p]$ and $R[u]$ are the sets of partial parallel execution trees reaching p and u, respectively.

These sets are then defined as the least solutions of the following constraint systems.

$$\begin{array}{llll} T[p] & \supseteq & T[r] & r \text{ return point of } p \quad (1) \\ T[s] & \supseteq & \{\#\} & s \text{ entry point of a procedure} \quad (2) \\ T[v] & \supseteq & e(T[u]) & e=(u,v) \text{ basic edge} \quad (3) \\ T[v] & \supseteq & e(T[u], T[p]) & e=(u,v) \text{ calls } p \quad (4) \\ T[v] & \supseteq & e(T[u], T[p_1] \,||\, T[p_2]) & e=(u,v) \text{ calls } p_1 \,||\, p_2 \quad (5) \end{array}$$

Line (1) says that the set of execution trees of procedure p is the set of execution trees for the return point of p. Line (2) says that at least the

empty tree is an execution tree for the entry point of a procedure. Line (3) says that for every basic edge $e = (u, v)$, the set of execution trees of program point v subsumes all execution trees to u extended by e. Line (4) says that for every edge $e = (u, v)$ calling a procedure p, the set of execution trees for the program point v subsumes all execution trees of u extended by e together with any execution tree of the procedure p. Finally, line (5) says that for every edge $e = (u, v)$ calling procedures p_1 and p_2 in parallel, the set of execution trees for the program point v subsumes all execution trees of u extended by e together with a pair of execution trees for p_1 and p_2, respectively.
Accordingly, the sets of reaching partial execution trees are given as the least solution of:

$$
\begin{array}{llll}
R[p] & \supseteq & R[u] & e = (u, _) \text{ calls } p,\ p \,||\, _ \text{ or } _ \,||\, p \qquad (6)\\
R[v] & \supseteq & R[v, \mathsf{main}] & \qquad (7)\\
R[v, p] & \supseteq & R_0[v] & v \text{ program point of } p \qquad (8)\\
R[v, p] & \supseteq & \bar{e}(R_0[u], R[v, q]) & e = (u, _) \text{ in } p \text{ calls } q \qquad (9)\\
R[v, p] & \supseteq & \bar{e}(R_0[u], R[v, q_1] \,||\, A[q_2]), & \\
 & & \bar{e}(R_0[u], A[q_1] \,||\, R[v, q_2]) & e = (u, _) \text{ in } p \text{ calls } q_1 \,||\, q_2 \qquad (10)\\
R_0[s] & \supseteq & \{\bullet\} & s \text{ entry node} \qquad (11)\\
R_0[v] & \supseteq & e(R_0[u]) & e = (u, v) \text{ basic edge} \qquad (12)\\
R_0[v] & \supseteq & e(R_0[u], T[p]) & e = (u, v) \text{ calls } p \qquad (13)\\
R_0[v] & \supseteq & e(R_0[u], T[p_1] \,||\, T[p_2]) & e = (u, v) \text{ calls } p_1 \,||\, p_2 \qquad (14)\\
A[p] & \supseteq & T[u] & u \text{ program point in } p \qquad (15)\\
A[p] & \supseteq & \bar{e}(T[u], A[q]) & e = (u, _) \text{ in } p \text{ calls } q \qquad (16)\\
A[p] & \supseteq & \bar{e}(T[u], A[q_1] \,||\, A[q_2]) & e = (u, _) \text{ in } p \text{ calls } q_1 \,||\, q_2 \qquad (17)
\end{array}
$$

Lines (6) and (7) reduce the sets of partial reaching execution trees to the sets $R[v, q]$ of partial execution trees reaching v relative to some procedure q. Lines (8) through (10) then explain how these sets of relative partial execution trees are determined. These constraints make use of two auxiliary groups of sets, namely, $R_0[v]$, v a program point, and $A[q]$, q a procedure. The sets $R_0[v]$ are defined in lines (11) through (14). They represent partial execution trees reaching v relative to the procedure of v. Finally, sets $A[q]$ being defined in lines (15) through (17), represent prefixes of computations running in parallel with the process currently under consideration.
By fixpoint induction, we verify:

PROPOSITION 5. *For all procedures p and program points u,*

$$
\begin{array}{llllll}
\Pi(p) & = & \mathcal{I}[T[p]] & \textit{and} & \Pi(u) & = & \mathcal{I}[T[u]]\\
\Pi_r(p) & = & \mathcal{I}[R[p]] & \textit{and} & \Pi_r(u) & = & \mathcal{I}[R[u]]
\end{array}
$$

□

By proposition 5, the semantics of programs based on execution paths can be understood as an abstraction of the semantics based on execution trees.

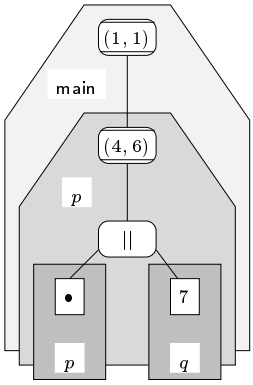

Fig. 4: The Process Tree Corresponding to Fig. 3.

Let us now turn to the concept of *process trees*. The internal nodes of process trees either denote entered but not yet completed calls e or splittings of the current process into sub-processes, whereas the sequence of leaves form a snapshot of all currently running processes.
The process tree corresponding to the partial execution tree from fig. 3 is shown in fig. 4. Note that within process trees, we now consider edges as *unary* tree symbols. Intuitively, the process tree corresponding to a partial execution tree is obtained by removing un-barred nodes. Just the barred nodes are retained since they correspond to the entered but not yet finished calls. Within these, the $||$-nodes combine trees of active processes running in parallel. Furthermore, the current program point of the process under focus is represented by $\bullet$, whereas the program points reached in parallel are recorded as leaves. In the example, the program point from q reached in parallel is 7, since so far no edge has been traversed inside q, and 7 is the entry point of q.
We now describe these concepts formally. *Process trees* for a parallel program can be represented as expressions c of the form:

$$\begin{array}{lcl} c & ::= & \bullet \mid e(c) \mid e(c \,||\, c') \mid e(c' \,||\, c) \\ c' & ::= & u \mid e(c') \mid e(c' \,||\, c') \end{array}$$

where $\bullet$ denotes the program point in the current process, e denotes a (possibly parallel) call edge, u a program point, and q a procedure called in parallel (here, we omitted barring of edges again).
Formally, we can define the extraction $\mathcal{C}[.]$ of process trees from partial

computation trees by:

$$
\begin{array}{lcll}
\mathcal{C}[\bullet] & = & \bullet & \\
\mathcal{C}[e(r)] & = & \mathcal{C}[r] & e \text{ basic edge} \\
\mathcal{C}[e(r,t)] & = & \mathcal{C}[r] & e \text{ call edge} \\
\mathcal{C}[\bar{e}(r,r_1)] & = & e(\mathcal{C}[r_1]) & e \text{ calls } q \\
\mathcal{C}[\bar{e}(r,r_1 \,||\, t)] & = & e(\mathcal{C}[r_1] \,||\, \mathcal{C}_q[t]) & e \text{ calls } _ \,||\, q \\
\mathcal{C}[\bar{e}(r,t \,||\, r_2)] & = & e(\mathcal{C}_q[t] \,||\, \mathcal{C}[r_2]) & e \text{ calls } q \,||\, _ \\
\mathcal{C}_q[\#] & = & u & u \text{ entry point of } q \\
\mathcal{C}_q[e(x)] & = & v & e = (_, v) \\
\mathcal{C}_q[e(x_1,x_2)] & = & v & e = (_, v) \\
\mathcal{C}_q[\bar{e}(x_1,x_2)] & = & \mathcal{C}[\bar{e}(x_1,x_2)] &
\end{array}
$$

The transformation $\mathcal{C}[.]$ can be seen as a (deterministic topdown) *finite-state linear tree transduction* from (partial) execution trees to process trees. Again, we extend it to a transformation of sets of (partial) trees by:

$$\mathcal{C}[T] \;=\; \{\mathcal{C}[t] \mid t \in T\}$$

For a program point u, let $C[u]$ denote the set of process trees possibly occurring at u (for brevity, we henceforth always state assertions just for program points u and omit the corresponding assertions for procedures p). The constraint system for $R[u]$ can be seen a regular tree grammar/finite state tree automaton where the constraint variables are the nonterminals/states. Let us call such a constraint system *regular*. Technically, a constraint system over sets of trees is regular, iff all right-hand sides of constraints are built up from constraint variables and constant sets of trees by application of tree constructors only. We conclude that the sets $R[u]$ are *regular tree languages*. As is well-known, regular tree languages are closed under application of finite-state linear tree transductions. Consequently, the sets $C[u]$ are regular as well.

8.2 Formalizing Refined Reachability Analysis

Consider again some abstract domain $\mathbb{D}$ together with a set $\mathbb{F}$ of transfer functions of the form $f\,x = (a \sqcap x) \sqcup b$, $a, b \in \mathbb{D}$ as in section 4. Assume we are given an analysis by some initial value $d_0 \in \mathbb{D}$ and a mapping $[.] : \mathsf{Basic} \to \mathbb{F}$ of basic edges to transfer functions. In the following, we will additionally assume that the lattice $\mathbb{D}$ is finite.
The basic **refined reachability** question consists in determining for all program points u and abstract values $d \in \mathbb{D}$, the sets

$$R[u,d] \;=\; \{t \in R[u] \mid \forall \pi \in \mathcal{I}[t] : [\pi]\, d_0 \sqsubseteq d\}$$

of all partial execution trees possibly reaching u such that *all* corresponding interleavings π result in values bounded by d. Indeed, given this information, we have perfect control over every possible history guaranteeing a certain

abstract value. The related analysis problem of Esparza et al. in [5, 6] is, in a certain sense, less ambitious than ours as it only asks for the sets

$$C[u,d] \quad = \quad \{\mathcal{C}[t] \mid t \in R[u], \forall \pi \in \mathcal{I}[t] : [\pi]\, d_0 \sqsubseteq d\}$$

of all process trees possibly reaching u such that *all* corresponding interleavings π result in values bounded by d. Also, it is only applicable to abstract lattices $\mathbb{D} = \{0 \sqsubset 1\}$. In order to relate the refined analysis problems for partial execution trees and process trees, respectively, we observe:

PROPOSITION 6. *For every program point u,* $\mathcal{C}[R[u,d]] = C[u,d]$. □

Due to this basic observation, we will first solve the refined analysis problem for partial execution trees, i.e., derive a regular constraint system for the sets $R[u,d]$. Regular constraint systems for the sets $C[u,d]$ then can be obtained just by application of the linear transduction $\mathcal{C}[.]$.
In order to construct a regular constraint system for the sets $R[u,d]$, we define an evaluation $\mathcal{H}[.]$ mapping (partial) execution trees to the least upper bound of the effects of all corresponding execution paths. The mapping $\mathcal{H}$ is given by:

$$\begin{array}{lcll}
\mathcal{H}[\#] & = & \mathcal{H}[\bullet] = \lambda x.x & \\
\mathcal{H}[e(x)] & = & [e] \circ \mathcal{H}[x] & e \text{ basic edge} \\
\mathcal{H}[e(x_1,x_2)] & = & \mathcal{H}[\bar{e}(x_1,x_2)] = \mathcal{H}[x_2] \circ \mathcal{H}[x_1] & e \text{ call edge} \\
H[x_1 \,||\, x_2] & = & \mathcal{H}[x_1] \otimes \mathcal{H}[x_2] &
\end{array}$$

Obviously, we have:

PROPOSITION 7. *For every (partial) execution tree x,* $\mathcal{H}[x] = [\mathcal{I}[x]]$. □

The evaluation function $\mathcal{H}[.]$ can be seen as a (bottom-up deterministic) finite tree automaton where the set of states is given by $\mathbb{F}$. Regular tree languages are closed under intersection. Therefore, $\mathcal{H}[.]$ together with the constraint system for reaching partial execution trees allows us to construct a regular constraint system for partial execution trees reaching program points in abstract states bounded by some d. This gives us our last main theorem:

THEOREM 4. *For every program point u and abstract value $d \in \mathbb{D}$, the sets $R[u,d]$ and $C[u,d]$ are regular. Regular constraint systems for these sets can be computed in time $\mathcal{O}(n \cdot m \cdot |\mathbb{F}|^3)$ where n is the size of the program and m the number of occurring procedures.* □

In particular, if $\mathbb{D} = \{0 \sqsubset 1\}$, then $|\mathbb{F}| = 3$, and we obtain at a quadratic algorithm. Note that we obtain explicit upper complexity bounds which moreover are tight in the sense that they match the sizes of the produced output structures.

9. Conclusion

We have shown how to extend the intra-procedural method of [11] to uniformly and efficiently capture inter-procedural bitvector analyses of fork/join parallel programs. Our method, which comprises analysis problems like available expressions, live variables or strong copy constant propagation, passes the test for practicality, as it 'behaves' as the widely accepted algorithms for sequential inter-procedural program analysis. Moreover, even though precision can only be proved for fork/join parallelism, our algorithm may also be used for computing safe approximations for languages with arbitrary synchronization statements. Also, due to its structural similarity to the sequential case, it can easily be integrated in program analysis environments like e.g. MetaFrame or PAG, which already contain the necessary fixpoint machinery.
Finally, we showed how to go beyond classical program analysis, which is characterized by taking the structure of the argument program dogmatically. In contrast, our refined analysis can be regarded as a generalization of the intra-procedural approach of property-oriented exansion to a setting with procedures and fork/join parallelism. Thus it does not focus on the structure of the argument program. Rather it provides information tailored to establish the appropriate program structure for the considered optimization goal. Technically, based on classical results for regular tree languages and specific properties of $\mathbb{F}$, we derived for every program point u and abstract value d, (a representation of) the sets of all partial execution trees as well as the sets of all process trees at u yielding abstract values bounded by d. We are convinced that the gained precision will be a basis for numerous agressive program optimization techniques.

References

[1] Martin Alt and Florian Martin. Generation of Efficient Interprocedural Analyzers with PAG. In *Proceedings of 2nd Static Analysis Symposium (SAS)*, pages 33–50. LNCS 983, Springer Verlag, 1995.

[2] Patrick Cousot. Semantic Foundations of Program Analysis. In Steven S. Muchnick and Neil D. Jones, editors, *Program Flow Analysis: Theory and Applications*, chapter 10, pages 303–342. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1981.

[3] Patrick Cousot and Radhia Cousot. Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints. In *Proceedings of 4th ACM Symposium on Principles of Programming Languages (POPL)*, pages 238–252. ACM Press, 1977.

[4] Patrick Cousot and Radhia Cousot. Static Determination of Dynamic Properties of Recursive Programs. In E.J. Neuhold, editor, *Formal Descriptions of Programming Concepts*, pages 237–277. North-Holland Publishing Company, 1977.

[5] J. Esparza and J. Knoop. An Automata-theoretic Approach to Interprocedural Dataflow Analysis. In *FoSSaCS '99*, volume 1578 of *Lecture Notes in Computer Science (LNCS)*, pages 14–30. Springer-Verlag, 1999.

[6] J. Esparza and A. Podelski. Efficient Algorithms for pre* and post* on Interprocedural Parallel Flow Graphs. In *ACM International Conference on Princples of Programming Languages (POPL)*, pages 1–11, 2000.

[7] M.S. Hecht. *Flow Analysis of Computer Programs.* The Computer Science Library. North-Holland, New York, 1977.
[8] J. Knoop. Parallel Constant Propagation. In *4th European Conference on Parallel Processing (Euro-Par)*, volume 1470 of *Lecture Notes in Computer Science (LNCS)*, pages 445–455. Springer-Verlag, 1998.
[9] J. Knoop, O. Rüthing, and B. Steffen. Towards a Tool Kit for the Automatic Generation of Interprocedural Data Flow Analyses. *Journal of Programming Languages*, 4(4):211–246, December 1996. Chapman & Hall, London (UK).
[10] J. Knoop and B. Steffen. The Interprocedural Coincidence Theorem. In *4th International Conference on Compiler Construction (CC'92)*, volume 641 of *Lecture Notes in Computer Science (LNCS)*, pages 125–140. Springer-Verlag, 1992.
[11] J. Knoop, B. Steffen, and J. Vollmer. Parallelism for Free: Efficient and Optimal Bitvector Analyses for Parallel Programs. *ACM Transactions on Programming Languages and Systems*, 18(3):268–299, 1996.
[12] E. Morel and C. Renvoise. Global Optimization by Suppression of Partial Redundancies. *Communications of the ACM (CACM)*, 22:96–103, 1979.
[13] H. Seidl and C. Fecht. Interprocedural Analyses: A Comparison. *Journal of Logic Programming (JLP)*, 43(2):123–156, 2000.
[14] B. Steffen. Property-Oriented Expansion. In *Proceedings of 2nd Static Analysis Symposium (SAS)*, pages 22–41. LNCS 1145, Springer Verlag, 1996.

Appendix A. Proof of Proposition 2

Consider functions $f_1, f_2, f \in \mathbb{F}$ given by $f_i = \lambda x.(a_i \sqcap x) \sqcup b_i$, $i = 1, 2$, and $f = \lambda x.(a \sqcap x) \sqcup b$. Then we calculate:

$$\begin{array}{rcl} f_1 \circ f_2 \sqcup f_2 \circ f_1 & = & (\lambda x.(a_1 \sqcap a_2 \sqcap x) \sqcup (a_1 \sqcap b_2) \sqcup b_1) \sqcup \\ & & (\lambda x.(a_1 \sqcap a_2 \sqcap x) \sqcup (a_2 \sqcap b_1) \sqcup b_2) \\ & = & \lambda x.(a_1 \sqcap a_2 \sqcap x) \sqcup b_1 \sqcup b_2 \\ & = & f_1 \otimes f_2 \end{array}$$

– giving us statement (1). Accordingly, for statement (2), we have:

$$\begin{array}{rcl} f_1 \otimes f \sqcup f_2 \otimes f & = & (\lambda x.(a_1 \sqcap a \sqcap x) \sqcup b_1 \sqcup b) \sqcup \\ & & (\lambda x.(a_2 \sqcap a \sqcap x) \sqcup b_2 \sqcup b) \\ & = & \lambda x.((a_1 \sqcup a_2) \sqcap a \sqcap x) \sqcup (b_1 \sqcup b_2) \sqcup b \\ & = & (f_1 \sqcup f_2) \otimes f \end{array}$$

Finally, for a proof of statement (3), let M_1, M_2 be non-empty subsets of E^*. By statement (1), we have

$$\begin{array}{rcll} [M_1] \otimes [M_2] & = & [M_1] \circ [M_2] \sqcup [M_2] \circ [M_1] & \\ & = & [M_2 \cdot M_1 \cup M_1 \cdot M_2] & \sqsubseteq \quad [M_1 \otimes M_2] \end{array}$$

Therefore, it remains to prove the reverse inequality. For that consider $w = e_1 \ldots e_m \in M_1 \otimes M_2$ where for disjoint index sets I_1, I_2 with $I_1 \cup I_2 = \{1, \ldots, m\}$, $w_i = w|_{I_i} \in M_i$. We claim:

$$[w] \sqsubseteq [w_1] \circ [w_2] \sqcup [w_2] \circ [w_1]$$

Clearly, this claim implies the statement (3) of our proposition. In order to prove the claim, let $[e_i] = \lambda x.(a_i \sqcap x) \sqcup b_i$, $i = 1, \ldots, m$, $[w_j] = \lambda x.(A_j \sqcap x) \sqcup B_j$, $j = 1, 2$, and $[w] = \lambda x.(A \sqcap x) \sqcup B$. Then by definition,

$$A = a_1 \sqcap \ldots \sqcap a_m = A_1 \sqcap A_2$$

Now consider value B. By definition,

$$B = \bigsqcup_{k=1}^{m} (b_k \sqcap a_{k+1} \sqcap \ldots \sqcap a_m)$$

We will show that for every k,

$$b_k \sqcap a_{k+1} \sqcap \ldots \sqcap a_m \sqsubseteq B_1 \sqcup B_2$$

W.l.o.g. assume that $k \in I_1$ (the case where $k \in I_2$ is completely analogous) and let $\{j_1, \ldots, j_r\} = \{j \in I_1 \mid j > k\}$. Then

$$b_k \sqcap a_{k+1} \sqcap \ldots \sqcap a_m \sqsubseteq b_k \sqcap a_{j_1} \sqcap \ldots \sqcap a_{j_r} \sqsubseteq B_1$$

which implies the assertion. □

Appendix B. Proof of Proposition 3

For $M \subset E^*$, let $E_M \subseteq E$ denote the set of elements of E occurring in M where $[e] = \lambda x.(a_e \sqcap x) \sqcup b_e$. Furthermore, let $B = \bigsqcup\{b_e \mid e \in E_M\}$.
We first show that $[\mathrm{pre}(M)] \sqsubseteq \lambda x.x \sqcup B$. Let $w \in \mathrm{pre}(M)$. We proceed by induction on the length n of w. If $n = 0$, then $w = \epsilon$. Consequently,

$$[w] = I = \lambda x.x \sqsubseteq \lambda x.x \sqcup B$$

Now assume $n > 0$. Then, $w = w'e$ where, by induction hypothesis, $[w']\, d \sqsubseteq d \sqcup B$ for every $d \in \mathbb{D}$. But then also

$$\begin{array}{rcll} [w]\, d & = & [e]\,([w']\, d) & \sqsubseteq \quad [e]\,(d \sqcup B) \\ & & & = \quad (a_e \sqcap d) \sqcup (a_e \sqcap B) \sqcup b_e \\ & & & \sqsubseteq \quad d \sqcup B \end{array}$$

which completes the proof. For a proof of the reverse inclusion, we first observe that $\lambda x.x = [\epsilon] \sqsubseteq [\mathrm{pre}(M)]$, since $\epsilon \in \mathrm{pre}(M)$. Furthermore, for each $e \in E_M$, we can find a prefix $w = w'e \in \mathrm{pre}(M)$. Consequently,

$$\begin{array}{rcl} b_e & \sqsubseteq & (a_e \sqcap [w']\, d) \sqcup b_e \\ & = & [w]\, d \sqsubseteq [\mathrm{pre}(M)]\, d \end{array}$$

for every $d \in \mathbb{D}$. Thus,

$$\lambda x.x \sqcup B = (\lambda x.x) \sqcup \bigsqcup\{\lambda x.b_e \mid e \in E_M\} \sqsubseteq [\mathrm{pre}(M)]$$

This completes the proof. □

Appendix C. Proof of Theorem 2

Let us start with the following simple but useful observation:

PROPOSITION 8. *For every $f \in \mathbb{F}$, $b \in \mathbb{D}$ and $\Delta = \lambda x.x \sqcup b$,*

$$f \otimes \Delta = \Delta \circ f \qquad \square$$

Next, we reformulate the constraint system for reachability as follows. We introduce the new values $[\![v]\!]'$, v a program point, and $[\![p]\!]'$, p a procedure, which collect the least upper bounds of *directly* reaching values by ignoring possible interleavings with execution paths possibly executed in parallel to v (or p). These values are determined as the least solution of the following constraint system:

$$\begin{array}{lll} [\![\mathsf{main}]\!]' \sqsupseteq d_0 & & (1) \\ [\![v]\!]' \sqsupseteq [v] \circ [\![p]\!]' & v \text{ program point in procedure } p & (2) \\ [\![p]\!]' \sqsupseteq [\![u]\!]' & e = (u, _) \text{ calls } p \text{ or } p \,||\, _ \text{ or } _ \,||\, p & (3) \end{array}$$

By standard fixpoint induction we find:

PROPOSITION 9. *For all program points v and procedures p,*

$$[\![v]\!] = [\![v]\!]' \sqcup B(v) \qquad \textit{and} \qquad [\![p]\!] = [\![p]\!]' \sqcup B(p) \qquad \square$$

In order to understand the "nature" of the values $B(v)$, we consider the sets $P(v)$ of edges possibly executed in parallel with program points v. They are determined through the least solution of the following constraint system:

$$\begin{array}{lll} E(p) \supseteq \{e\} & e \text{ basic edge in procedure } p \\ E(p) \supseteq E(q) & \text{procedure } p \text{ calls } q \text{ or } q \,||\, _ \text{ or } _ \,||\, q \\ P(v) \supseteq P(p) & v \text{ program point in } p \\ P(p) \supseteq P(u) & (u, _) \text{ calls procedure } p \\ P(q_i) \supseteq E(q_{3-i}) \cup P(u) & (u, _) \text{ calls } q_1 \,||\, q_2 \end{array}$$

By comparison of this constraint system with the definition of the values $\sigma(p)$ and $B(v), B(p)$ in section 6, we obtain:

PROPOSITION 10. *For every procedure p and program point v,*

(1) $\lambda x.x \sqcup \sigma(p) = I \sqcup [E(p)]$;

(2) $\lambda x.x \sqcup B(p) = I \sqcup [P(p)]$ *and* $\lambda x.x \sqcup B(v) = I \sqcup [P(v)]$. $\square$

Moreover, we have:

PROPOSITION 11. *For every procedure p,*

$$[pre(\Pi(p))] = I \sqcup \bigsqcup [E(p)] = \lambda x.x \sqcup \sigma(p) \qquad \square$$

In order to simplify the proof of theorem 2, let us assume that all calls are parallel calls $q_1 \,||\, q_2$. This assumption does not incur a restriction, since an ordinary call to a procedure p can easily be simulated by a call to $p \,||\, q_0$

where q_0 is a procedure with just a single program point and no edges at all. Furthermore, it suffices to prove the assertion of the theorem just for program points v (the assertion for procedures then is an immediate consequence). We want to prove that for every program point v, the value $[\![v]\!]$ is a safe approximation of the value $\mathsf{Reach}(v)$, i.e., $[\![v]\!] \sqsupseteq \mathsf{Reach}(v)$. By definition,

$$\mathsf{Reach}(v) = [\Pi_r(v)]\, d_0 = [\Pi(v, \mathsf{main})]\, d_0$$

Therefore, let $w \in \Pi(v, \mathsf{main})$. Then there are program points $u_0, \ldots, u_m$, execution paths $w_0, \ldots, w_m$ together with execution paths w_i', procedures $q_1^{(i)}, q_2^{(i)}$ and indices $j(i) \in \{1, 2\}$ for $i = 1, \ldots, m$ such that:

- $u_m = v$;
- $w_i \in \Pi(u_i)$ for $i = 0, \ldots, m$;
- there are calls $(u_{i-1}, _)$ to $q_1^{(i)} \,||\, q_2^{(i)}$;
- u_0 is a program point in main and for $i > 0$, u_i is a program point in $q_{j(i)}^{(i)}$;
- $w_i' \in \mathrm{pre}(\Pi(q_{3-j(i)}^{(i)}))$ for $i = 1, \ldots, m$;
- $w \in \{w_0\}\cdot(\{w_1'\}\otimes(\{w_1\}\cdot(\ldots\{w_{m-1}'\}\otimes(\{w_{m-1}\}\cdot(\{w_m'\}\otimes\{w_m\}))\ldots)))$.

Let $\Delta = \lambda x.x \sqcup P(v)$. Then by proposition 11,

$$[w_i'] \sqsubseteq [\mathrm{pre}(\Pi(q_{3-j(i)}^{(i)}))] = I \sqcup [E(q_{3-j(i)}^{(i)})] \sqsubseteq I \sqcup [P(v)] = \Delta$$

for all $i = 1, \ldots, m$. Therefore by proposition 8,

$$\begin{array}{rcl}
[w] & \sqsubseteq & (((\ldots(([w_m] \otimes [w_m']) \circ [w_{m-1}]) \otimes [w_{m-1}']\ldots) \circ [w_1]) \otimes [w_1']) \circ [w_0] \\
& \sqsubseteq & (((\ldots(([w_m] \otimes \Delta) \circ [w_{m-1}]) \otimes \Delta\ldots) \circ [w_1]) \otimes \Delta) \circ [w_0] \\
& = & \Delta \circ (\ldots(([w_m] \otimes \Delta) \circ [w_{m-1}]) \otimes \Delta\ldots) \circ [w_1] \circ [w_0] \\
& \ldots & \\
& = & \Delta \circ [w_m] \circ [w_{m-1}] \circ \ldots \circ [w_0]
\end{array}$$

Since $([w_m] \circ \ldots \circ [w_0])\, d_0 \sqsubseteq [\![v]\!]'$, we conclude that

$$[w]\, d_0 \sqsubseteq \Delta\, [\![v]\!]' = [\![v]\!]' \sqcup B(v) = [\![v]\!]$$

which we wanted to prove.

It remains to prove the reverse inequality, i.e., that (1) $[\![v]\!]' \sqsubseteq \mathsf{Reach}(v)$ and (2) $B(v) \sqsubseteq \mathsf{Reach}(v)$.

Let us first consider inequality (1). The value $[\![v]\!]'$ is the least upper bound on values $[w]\, d_0$ such that there exit program points $u_0, \ldots, u_m$, execution paths $w_0, \ldots, w_m$ together with procedures $q_1^{(i)}, q_2^{(i)}$ and indices $j(i) \in \{1, 2\}$ for $i = 1, \ldots, m$ such that:

- $u_m = v$;
- $w_i \in \Pi(u_i)$ for $i = 0, \ldots, m$;

- there are calls $(u_{i-1}, _)$ to $q_1^{(i)} \,||\, q_2^{(i)}$;
- u_0 is a program point in main and for $i > 0$, u_i is a program point in $q_{j(i)}^{(i)}$;
- $w = w_0 \dots w_m$.

By induction on $r = m - i$ (from $r = 0$ to $r = m - 1$), we find that for $i > 0$,

$$w_i \dots w_m \in \Pi(v, q_{j(i)}^{(i)})$$

and for $i = 0$,

$$w = w_0 \dots w_m \in \Pi(v, \mathsf{main}) = \Pi_r(v)$$

Therefore,

$$[w]\, d_0 \sqsubseteq [\Pi_r(v)]\, d_0 = \mathsf{Reach}(v)$$

which we wanted to prove.
Now let us consider inequality (2). By proposition 10, $\lambda x.x \sqcup B(v) = I \sqcup [P(v)]$. Therefore, it suffices to prove for each edge $e \in P(v)$, that $b_e \sqsubseteq \mathsf{Reach}(v)$.
Since $e \in P(v)$, there exist program points $u_0, \dots, u_m$, execution paths $w_0, \dots, w_m$ together with procedures $q_1^{(i)}, q_2^{(i)}$, indices $j(i) \in \{1, 2\}$ for $i = 1, \dots, m$, an index $k \in \{1, \dots, m\}$ and one execution path w' such that

- $u_m = v$;
- $w_i \in \Pi(u_i)$ for $i = 0, \dots, m$;
- there are calls $(u_{i-1}, _)$ to $q_1^{(i)} \,||\, q_2^{(i)}$;
- u_0 is a program point in main and for $i > 0$, u_i is a program point in $q_{j(i)}^{(i)}$;
- $w'e \in \mathrm{pre}(\Pi(q_{3-j(k)}^{(k)})$.

As above, we conclude that $w_k \dots w_m \in \Pi(v, q_{j(k)}^{(k)})$. By definition, then also

$$w_{k-1} w_k \dots w_m w' e \in \Pi(v, q_{j(k-1)}^{(k-1)})$$

(where in case $k = 1$, we let $q_{j(0)}^{(0)} = \mathsf{main}$) and therefore also

$$w_0 \dots w_{k-1} w_k \dots w_m w' e \in \Pi(v, \mathsf{main}) = \Pi_r(v)$$

We conclude that

$$b_e \sqsubseteq b_e \sqcup (a_e \sqcap ([w_0 \dots w_m w']\, d_0)) = [w_0 \dots w_m w' e]\, d_0 \sqsubseteq [\Pi_r(v)]\, d_0 = \mathsf{Reach}(v)$$

which completes the proof. $\square$