

INTRODUCTION

partial evaluation

- Given a program and **part** of its input data—the so called **static** data—a partial evaluator returns a new, **residual** program which is specialized for the given data

Motivation

- Very few approaches devoted to formally analyze the effects of partial evaluation (mainly experimental)

We introduce a **symbolic** approach for predicting the potential effects of PE:

- generate a finite representation that safely describes all possible **call traces**
- analyze how this finite representation would **change** by a particular partial evaluation
- compare the original and the transformed representations
(in some cases one can **predict** the effects of running the partial evaluator beforehand)

TRACE ANALYSIS FOR LOGIC PROGRAMS

Overview of the method

- construct a context-free grammar (CFG) that approximates the call traces of a logic program (LP)
- approximate the CFG by a strongly regular grammar (SRG) (if needed)
- transform the SRG into a finite automaton (FA)

$$\text{LP} \Rightarrow \text{CFG} \Rightarrow \text{SRG} \Rightarrow \text{FA}$$

Definition (call trace)

We say that $\tau = p_0 p_1 \dots p_{n-1} \in \Pi^*$, $n \geq 1$, is a call trace for Q_0 with P iff

there exists a successful SLD derivation $Q_0 \xrightarrow{p_0} Q_1 \xrightarrow{p_1} \dots \xrightarrow{p_{n-1}} Q_n$

where each SLD step is labeled with the predicate symbol of the selected atom

FROM LOGIC PROGRAMS TO CONTEXT-FREE GRAMMARS

- A CFG is a tuple $G = \langle \Sigma, N, R, S \rangle$, where Σ and N are two disjoint sets of **terminals** and **non-terminals**, respectively, $S \in N$ is the **start** symbol, and R is a set of **rules**

Definition (trace CFG)

Let P be a program and $q \in \Pi$ a predicate symbol. The associated trace CFG is $\text{CFG}_q^P = \langle \Pi, \bar{\Pi} \cup \{\text{START}\}, R, \text{START} \rangle$, where

$$R = \{ \text{START} \rightarrow \bar{q} \} \cup \{ \overline{\text{pred}(A_0)} \rightarrow \overline{\text{pred}(A_0)} \overline{\text{pred}(B_1)} \dots \overline{\text{pred}(B_n)} \mid A_0 \leftarrow B_1, \dots, B_n \in P, n \geq 0 \}$$

where

- $\overline{\text{pred}(A)}$ returns a **non-terminal** associated to the predicate symbol of A
- $\text{pred}(A)$ returns a **terminal** associated to the predicate symbol of A

CFG_q^P is a **correct approximation** of the call traces for P w.r.t. the leftmost computation rule

EXAMPLE

- (c₁) $\text{mlist}([], I, [])$.
(c₂) $\text{mlist}([X|R], I, L) \leftarrow \text{ml}(X, R, I, L)$.
(c₃) $\text{ml}(X, R, I, [X|R]) \leftarrow \text{mult}(X, I, XI), \text{mlist}(R, I, RI)$.
(c₄) $\text{mult}(0, Y, 0)$. (c₅) $\text{mult}(s(X), Y, Z) \leftarrow \text{mult}(X, Y, Z1), \text{add}(Z1, Y, Z)$.
(c₆) $\text{add}(X, 0, X)$. (c₇) $\text{add}(X, s(Y), s(Z)) \leftarrow \text{add}(X, Y, Z)$.

Associated trace CFG:

$$\text{CFG}_{\text{mlist}}^P = \{ \text{mlist}, \text{ml}, \text{mult}, \text{add} \}, \{ \text{START}, \text{MLIST}, \text{ML}, \text{MULT}, \text{ADD} \}, R, \text{START}$$

where the set of rules R is as follows:

$$\begin{array}{lll} \text{START} \rightarrow \text{MLIST} & \text{ML} \rightarrow \text{ml MULT MLIST} & \\ \text{MLIST} \rightarrow \text{mlist} & \text{MULT} \rightarrow \text{mult} & \text{ADD} \rightarrow \text{add} \\ \text{MLIST} \rightarrow \text{mlist ML} & \text{MULT} \rightarrow \text{mult MULT ADD} & \text{ADD} \rightarrow \text{add ADD} \end{array}$$

FROM CONTEXT-FREE GRAMMARS TO STRONGLY REGULAR GRAMMARS

Basic idea

- We use Mohri and Nederhof's technique for approximating CFGs with SRGs, which can then be mapped to equivalent finite-state automata
- We consider **left-linear** grammars: every rule has either the form $(A \rightarrow t)$ or $(A \rightarrow tB)$ where t is a finite sequence of terminals and A, B are non-terminals

Definition (trace SRG)

For each set M of mutually recursive non-terminals such that their rules are not all left-linear w.r.t. the non-terminals of M , we apply the following transformation:

- For each non-terminal $A \in M$, we introduce a fresh non-terminal A' and add the rule $A' \rightarrow \epsilon$
- For each non-terminal $A \in M$ and each rule

$$A \rightarrow t_0 B_1 t_1 B_2 t_2 \dots B_m t_m \text{ of } \text{CFG}_q^P$$

with $m \geq 0, B_1, \dots, B_m \in M, t_0, \dots, t_m \in (\Pi \cup (\bar{\Pi} \setminus M))^*$, we replace this rule by the following set:

$$\begin{array}{l} A \rightarrow t_0 B_1 \\ B'_1 \rightarrow t_1 B_1 \\ \dots \\ B'_{m-1} \rightarrow t_{m-1} B_m \\ B'_m \rightarrow t_m A' \end{array}$$

- We let $\text{SRG}_q^P = \langle \Pi, \bar{\Pi} \cup N \cup \text{START}, R', \text{START} \rangle$, where R' are the rules obtained as described above and N are the fresh non-terminals added during this process.

EXAMPLE

- The sets of mutually recursive non-terminals are

$$\{ \{ \text{MLIST}, \text{ML} \}, \{ \text{MULT} \}, \{ \text{ADD} \} \}$$

- Therefore,

- The rules for **MLIST** and **ML** are left-linear w.r.t. $\{ \text{MLIST}, \text{ML} \}$
- The rules for **ADD** are clearly left-linear too

- However, the second rule of **MULT**:

$$\text{MULT} \rightarrow \text{mult MULT ADD}$$

- is not left-linear because, even if **ADD** is treated as a terminal, it appears **to the right** of the non-terminal **MULT**

- Therefore, in $\text{SRG}_{\text{mlist}}^P$ we replace the original rules for **MULT** by the following ones:

$$\begin{array}{l} \text{MULT}' \rightarrow \epsilon \\ \text{MULT} \rightarrow \text{mult MULT}' \\ \text{MULT} \rightarrow \text{mult MULT} \\ \text{MULT}' \rightarrow \text{ADD MULT}' \end{array}$$

From SRGs to Finite Automata (and Regular Expressions)

- The language associated to a SRG can be represented by
 - a finite-state automaton (FA) or
 - a regular expression (RE)

Trace FA

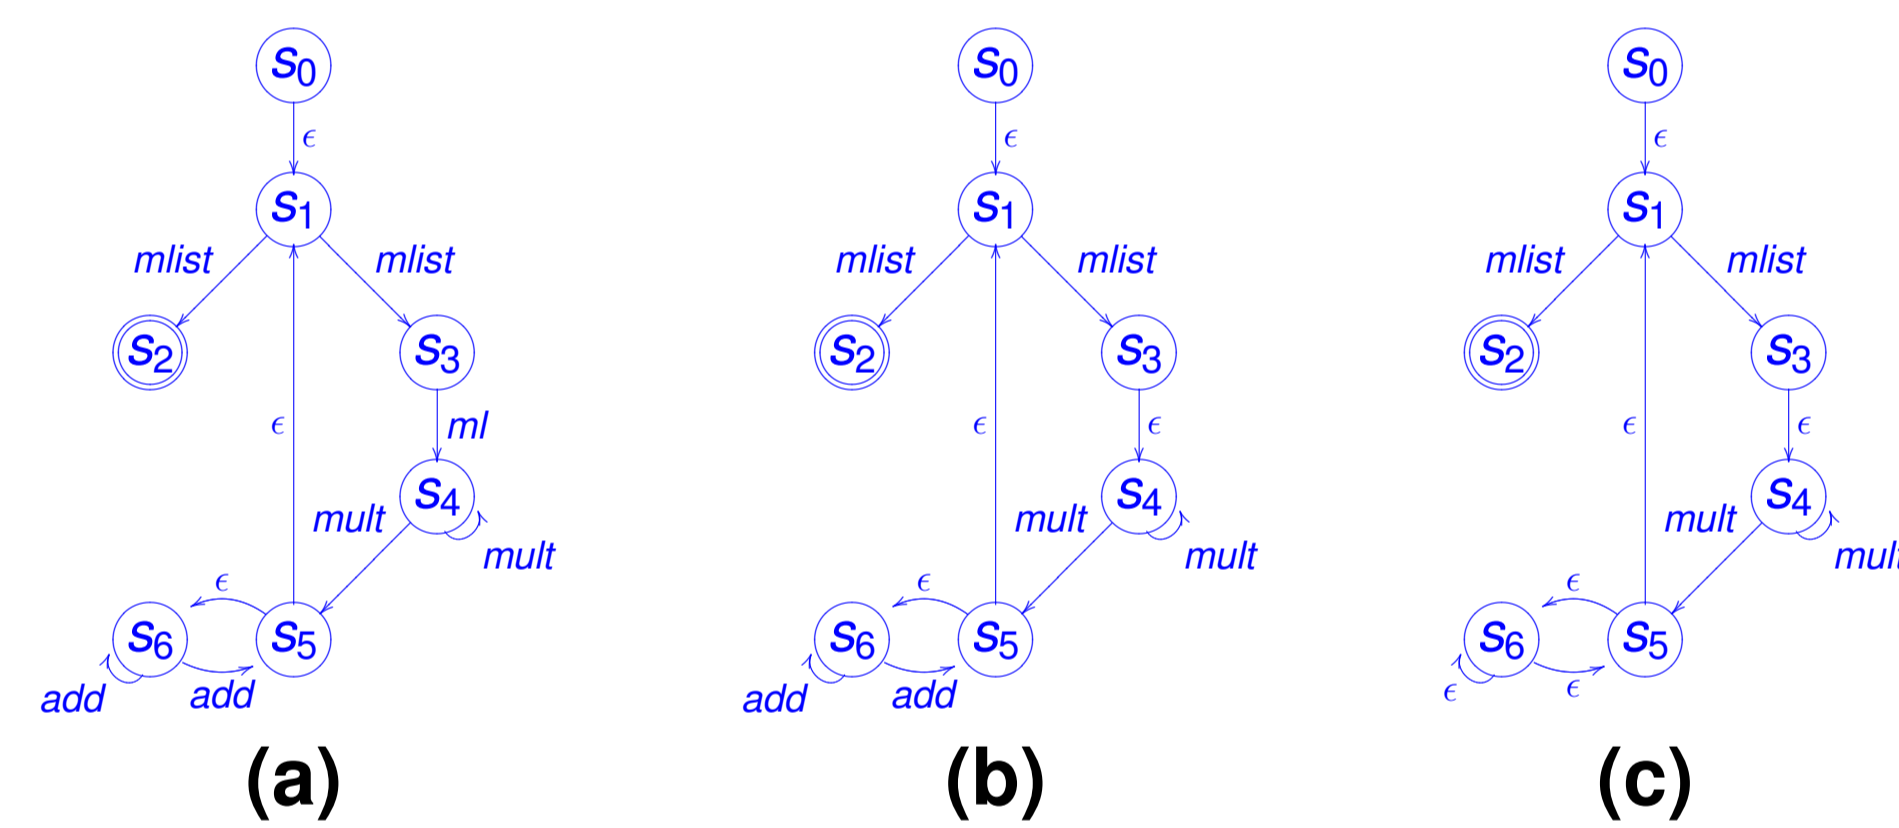
- A FA is specified by a tuple $\langle Q, \Sigma, \delta, s_0, F \rangle$, where
 - Q is a set of states
 - Σ is an input alphabet
 - $\delta \subseteq Q \times \Sigma \times Q$ is a set of transitions
 - $s_0 \in Q$ is the start state
 - $F \subseteq Q$ is a set of final states

- We follow the classical approach to construct a FA from a SRG

- there is a start state associated to the start symbol of the SRG
- for each reduction $w \rightarrow w'$ with a rule $A \rightarrow tB$, we have a transition (s, α, s') in the FA, where states s, s' are associated with the sequence of non-terminals in w, w' and character α is set to the sequence t in the applied rule

EXAMPLE

The trace FA associated to $\text{SRG}_{\text{mlist}}^P$ is the FA (a) below:



where

$$\begin{array}{lll} S_0 = \text{START} & S_1 = \text{MLIST} & S_2 = \epsilon \quad \text{final state} \\ S_3 = \text{ML} & S_4 = \text{MULT MLIST} & S_5 = \text{MULT}' \text{ MLIST} \\ S_6 = \text{ADD MULT}' \text{ MLIST} & & \end{array}$$

TOWARDS PREDICTING THE SPEEDUP OF PARTIAL EVALUATION

- The trace analysis gives us the **context** where every predicate call appears
- We introduce two **transformations** that modify the computed traces to account for the potential effects of a partial evaluation

Elimination of intermediate predicates

- For every state with exactly one input transition and one output transition, replace the label of the output transition by ϵ (i.e., delete calls to predicates which are called from a single program point)
- For the trace FA (a) above, we get the trace FA (b) by eliminating the intermediate state s_3

Removal of unfoldable predicates

- The transformation is parameterized by the output of a BTA, which annotates each predicate with either **unfold** or **memo**. Basically, the labels of unfoldable predicates are replaced with ϵ
- Given a BTA that annotates *mlist*, *ml*, and *mult* with **memo** and *add* with **unfold**, the trace FA (b) above is transformed into the trace FA (c)
- Is this PE useful?** YES, we will achieve a significant improvement since, in every iteration for *mlist*, we will save the (recursive) evaluation of the calls to *add*

CONCLUSIONS

- The closest approach to our trace analysis is that of Gallagher and Lafave (1996), though we offer a different trade-off between analysis cost and accuracy:
 - they generate trace **terms** abstracting computation trees **independently of a computation rule**, while we generate **sequences** of predicate calls for a **specific computation rule**
 - they do not include a technique for enumerating the (possibly infinite) set of trace terms of a program, while this is a key ingredient of our approach
- A proof-of-concept implementation of our technique, called PEPE, is publicly available from <http://german.dsic.upv.es/pepe.html>
- Our approach is a first step towards the development of automated techniques and tools for predicting the potential speedup of partial evaluation