

Understanding the Origin of Alarms in ASTRÉE

Xavier Rival

École Normale Supérieure, 45, rue d'Ulm,
75230, Paris cedex 5, France

Abstract. Static analyzers like ASTRÉE are incomplete, hence, may produce false alarms. We propose a framework for the investigation of the alarms produced by ASTRÉE, so as to help classifying them as *true errors* or *false alarms* that are due to the approximation inherent in the static analysis. Our approach is based on the computation of an approximation of a set of traces specified by an initial and a (set of) final state(s). Moreover, we allow for finer analyses to focus on some execution patterns or on some possible inputs. The underlying algorithms were implemented inside ASTRÉE and used successfully to track alarms in large, critical embedded applications.

1 Introduction

The risk of failure due to software bugs is no longer considered acceptable in the case of critical applications (as in aerospace, energy, automotive systems). Therefore, sound program analyzers have been developed in the last few years, that aim at proving safety properties of critical, embedded software such as memory properties [22], absence of runtime errors [3], absence of buffer overruns [11], correctness of pointer operations [28]. These tools attempt to prove automatically the correctness of programs, even though this is not decidable; they are sound (they never claim the property of interest to hold even though it does not) and always terminate; hence, they are *incomplete*: they may produce false alarms, i.e. report not being able to prove the correctness of some critical operation even though no concrete execution of the program fails at this point.

Alarms are a major issue for end-users. Indeed, in case the analyzer reports an alarm, it could either be a false alarm or a real bug that should be fixed. Ideally, a report for a true error should come with an error scenario. Currently, the alarm investigation process in ASTRÉE [3] mainly relies on the manual inspection of invariants, partly with a graphical interface [10]; this process turns out to be cumbersome, since even simple alarms may take days to classify.

A false alarm might either be due to an imprecision of some abstract operations involved in the analysis (e.g. the abstract join operator usually loses precision) as in Fig. 1(a) (simplified version of an alarm formerly reported by ASTRÉE) or to a lack of expressiveness of the domain (checking the example of Fig. 1(c) requires proving a very deep arithmetic theorem). In the former case, we may expect a (semi)-automatic refinement process to prove the alarm to be

<pre> l₀ : input(x); l₁ : if(x > 0){y = x; } else{y = -x; } l₂ : b = (y > 10); l₃ : assert(b ⇒ (x < -10 ∨ 10 < x)); (a) </pre>	<pre> l₀ : x = 1; y = 1; while(true){ l_i : input(x); l_d : assert(y > 0); y = x; } (b) </pre>	<pre> x, y, z are integer variables input(x); input(y); input(z); if(x > 0 ∧ y > 0 ∧ z > 0){ assert(x⁴ ≠ y⁴ + z⁴); } (c) </pre>
--	--	--

Fig. 1. Three example programs

false; in the latter, the design of a refined domain can hardly be automated, so we can only hope for a refined alarm description.

Our goal is to provide some support in the alarm investigation process. We propose to resort to automatic, sound static analysis techniques so as to refine an initial static analysis into an approximation of a subset of traces that actually lead to an error. If a combination of forward and backward refining analyses allows to prove that this set is empty, we can conclude the alarm is false (as in Fig. 1(a)); otherwise, we get a refined characterization of the (possibly fictitious) erroneous traces. We propose to refine this kind of *semantic slicing* (i.e. extraction of part of the program traces) by selecting some alarm contexts (e.g. traces leading to an error after two iterations in a loop or traces constrained by some set of inputs). A similar process allows to check an error scenario, by slicing the traces reaching the alarm point in some context specified by an execution pattern and a set of inputs that are supposed to be valid: in case the analysis reveals that such conditions *always* lead to an error, the error scenario is validated (this can be achieved in the example of Fig. 1(b)); it is a counter-example.

The contribution of the paper is both theoretical and practical:

- we propose a framework for alarm inspection, based on backward analysis [6,7], trace partitioning [23], and slicing techniques [29,19];
- we provide encouraging practical results obtained by an implementation inside the ASTRÉE static analyzer [3,10].

Sect. 2 introduces forward-backward analysis-based approximation (semantic slicing) of a set of traces resulting in an error. Semantic slicing refinements are introduced in Sect. 3: restriction to some execution patterns in Sect. 3.1 and to some inputs in Sect. 3.2. Sect. 4 applies syntactic slicing techniques to the reduction of the amount of code to analyze. Sect. 5 presents some practical examples. Sect. 6 concludes and reviews related work.

2 Backward Analysis

2.1 Standard Notations

We restrict to a subset of \mathbb{C} for the sake of concision. We let \mathbb{X} (resp. \mathbb{V}) denote the set of variables (resp. of values); we write \mathfrak{e} (resp. \mathfrak{s}) for the set of expressions (resp. statements aka programs). Variables and values have an integer, floating-point or boolean type. We consider assignments, conditionals, loops, assertions, and an additional **input**(x) statement, that emulates the reading of a volatile

variable x : this statement replaces the value of x with a random value of the corresponding type. The grammar is given below. The control point before each statement and at the end of each block is associated to a unique label $l \in \mathbb{L}$.

$$\begin{aligned}
 e \ (e \in \mathfrak{E}) &::= v \ (\text{where } v \in \mathbb{V}) \mid x \ (\text{where } x \in \mathbb{X}) \mid e \oplus e \\
 s \ (s \in \mathfrak{S}) &::= x := e \ (\text{where } x \in \mathbb{X}, e \in \mathfrak{E}) \mid s; s \mid \mathbf{skip} \\
 &\mid \mathbf{input}(x) \mid \mathbf{assert}(e) \mid \mathbf{if}(e)\{s\}\mathbf{else}\{s\} \mid \mathbf{while}(e)\{s\}
 \end{aligned}$$

In practice, the subset of \mathbb{C} we analyze also includes functions, pointers, composite data-structures, all kinds of definitions, and all control structures. It excludes recursive functions, dynamic memory allocation, and destructive updates.

A state $s \in \mathbb{S}$ is either the error state Ω or a pair (l, ρ) where l is a label and $\rho \in \mathbb{M} = \mathbb{X} \rightarrow \mathbb{V}$ is a *memory state* (aka *store*). Note that we assume there are no errors in expressions; hence, the error state Ω can only be reached after a failed assertion. The semantics $\llbracket s \rrbracket = \langle \sigma_0, \dots, \sigma_n \rangle \mid \forall i, \sigma_i \rightarrow \sigma_{i+1}$ of a program s is a set of sequences of states (aka traces), such that any two successive states are related by the transition relation \rightarrow of the program. The relation \rightarrow is defined by local rules, such as the following:

- assert statement $l_0 : \mathbf{assert}(e); l_1$: if $\llbracket e \rrbracket(\rho) = \mathbf{true}$ ($\llbracket e \rrbracket(\rho)$ is the result of the evaluation of e in ρ), then $(l_0, \rho) \rightarrow (l_1, \rho)$; if $\llbracket e \rrbracket(\rho) = \mathbf{false}$, then $(l_0, \rho) \rightarrow \Omega$.
- assignment $l_0 : x := e; l_1$: $(l_0, \rho) \rightarrow (l_1, \rho[x \leftarrow \llbracket e \rrbracket(\rho)])$ (where $\llbracket e \rrbracket \in \mathbb{M} \rightarrow \mathbb{V}$);
- input statement $l_0 : \mathbf{input}(x); l_1$: if $v \in \mathbb{V}$ has the same type as x , then $(l_0, \rho) \rightarrow (l_1, \rho[x \leftarrow v])$;

2.2 Approximation of Dangerous Traces

Dangerous States. We consider a program $P \in \mathfrak{s}$. A state σ is *dangerous* if it is not Ω and may lead to Ω in one transition step: $\sigma \rightarrow \Omega$. A *dangerous label* is a label l followed by an assertion statement ($l : \mathbf{assert}(e);$). ASTRÉE over-approximates the set of *reachable* dangerous states; hence, our goal is to start with such an over-approximation and to make a diagnosis whether a set of concrete, dangerous states is actually reachable.

Dangerous Traces. First, we are interested in real executions only, that is traces starting from an *initial* state. We let $\mathcal{I} \subseteq \mathbb{S}$ denote the set of initial states. Then, the set of real executions is $\vec{T} = \langle \sigma_0, \dots, \sigma_n \rangle \in \llbracket P \rrbracket \mid \sigma_0 \in \mathcal{I} \rangle = \mathbf{lfp}_\emptyset \vec{F}$ where $\vec{F} : E \mapsto \{ \langle \sigma \rangle \mid \sigma \in \mathcal{I} \} \cup \{ \langle \sigma_0, \dots, \sigma_n, \sigma_{n+1} \rangle \mid \langle \sigma_0, \dots, \sigma_n \rangle \in E \wedge \sigma_n \rightarrow \sigma_{n+1} \}$ constructs execution traces forward, and $\mathbf{lfp}_X F$ is the least fixpoint of F greater than X . Second, we restrict to executions ending in a dangerous state (or at a dangerous label). We let $\mathcal{F} \subseteq \mathbb{S}$ denote the set of final states of interest. The set of executions ending in \mathcal{F} is $\overleftarrow{T} = \langle \sigma_0, \dots, \sigma_n \rangle \in \llbracket P \rrbracket \mid \sigma_n \in \mathcal{F} \rangle = \mathbf{lfp}_\emptyset \overleftarrow{F}$ where \overleftarrow{F} is defined in a similar way as \vec{F} : $\overleftarrow{F} : E \mapsto \{ \langle \sigma \rangle \mid \sigma \in \mathcal{F} \} \cup \{ \langle \sigma_{-1}, \sigma_0, \dots, \sigma_n \rangle \mid \langle \sigma_0, \dots, \sigma_n \rangle \in E \wedge \sigma_{-1} \rightarrow \sigma_0 \}$.

The set of traces of interest is $T = \vec{T} \cap \overleftarrow{T} = \mathbf{lfp}_\emptyset \vec{F} \cap \mathbf{lfp}_\emptyset \overleftarrow{F}$. It is a subset of all program behaviors $\llbracket P \rrbracket$; in this respect, we call T a *semantic slice*.

In the following, \mathcal{F} may represent either $\mathcal{F}_l = \{(l, \rho) \mid \rho \in \mathbb{M}\}$ or $\mathcal{F}_{\mathfrak{D},l} = \{(l, \rho) \mid \rho \in \mathbb{M} \wedge (l, \rho) \rightarrow \Omega\}$, unless we specify explicitly; the slice T_l (resp. $T_{\mathfrak{D},l}$) defined by \mathcal{F}_l (resp. $\mathcal{F}_{\mathfrak{D},l}$) collects the executions ending at label l (resp. the executions causing an error at label l). T shall represent either T_l or $T_{\mathfrak{D},l}$.

Example 1. In the code of Fig. 1(a), label l_3 is a dangerous label; the set of dangerous states for the corresponding **assert** is $\mathcal{F}_{\mathfrak{D},l_3} = \{(l_3, \rho) \mid \rho \in \mathbb{M} \wedge \rho(b) = \mathbf{true} \wedge -10 \leq \rho(x) \leq 10\}$. The set of initial states is $\mathcal{I} = \{(l_0, \rho) \mid \rho \in \mathbb{M}\}$. Clearly, this program does not cause any error: if $y > 10$ at l_2 , then, either $x > 0$ and $x = y > 10$ or $x \leq 0$ and $x = -y < -10$. Hence, we wish to prove that $T_{\mathfrak{D},l_3} = \emptyset$.

Alarm Inspection. Our goal is to determine whether an alarm is true or not. We may fall in either of the following cases:

Case a) **alarm proved false:** if the static analysis proves that $T_{\mathfrak{D},l} = \emptyset$, then the dangerous states in $\mathcal{F}_{\mathfrak{D},l}$ are not reachable and *the alarm is false*;

Case b) **alarm proved true:** if the static analysis proves that any trace in T_l violates the assert (i.e. all traces reaching l cause an error at this point), then *the alarm is a true error*;

Case c) **undecided case:** obviously, we may not be able to conclude to either of the previous cases; then, either *an error would occur in some cases only* (this case is considered in Sect. 3) or *the lack of conclusion* is due to a lack of expressivity of the abstract domain (the refined analysis of the alarm context should help designing a domain refinement).

Trace Approximation. The approximation of sets of traces is based on an abstraction of sets of stores defined by a domain $(D_n^\sharp, \sqsubseteq)$ and a concretization function $\gamma_n : D_n^\sharp \rightarrow \mathcal{P}(\mathbb{M})$ [8]. We assume that D_n^\sharp provides a widening operator ∇ , approximate binary lub (\sqcup) and glb (\sqcap) operators, and a least element \perp , with the usual soundness assumptions. The domain for approximating sets of executions is defined by $D^\sharp = \mathbb{L} \rightarrow D_n^\sharp$ and $\gamma : (I \in D^\sharp) \mapsto \{\langle (l_0, \rho_0), \dots, (l_n, \rho_n) \rangle \mid \forall i, \rho_i \in \gamma_n(I(l_i))\}$: a set of traces is approximated by local invariants, approximating the sets of stores that can be encountered at any label. We let \mathbf{lfp}^\sharp denote an abstract post-fixpoint operator, derived from ∇ : if $F : D \rightarrow D$, $F \circ \gamma \subseteq \gamma \circ F^\sharp$ and $X \subseteq \gamma(X^\sharp)$, then $\mathbf{lfp}_X F \subseteq \gamma(\mathbf{lfp}_X^\sharp F^\sharp)$. The domain D_n^\sharp is supposed to feature sound abstract operations $\overline{\text{assign}} : \mathbb{X} \times \mathbb{e} \times D_n^\sharp \rightarrow D_n^\sharp$, $\text{guard} : \mathbb{e} \times D_n^\sharp \rightarrow D_n^\sharp$, $\text{forget} : \mathbb{X} \times D_n^\sharp \rightarrow D_n^\sharp$ that soundly mimic the concrete assignment, testing of conditions, and reading of inputs (by forgetting the value of the modified variable). For instance, the soundness of *guard* boils down to $\forall d \in D_n^\sharp, \forall \rho \in \gamma_n(d), \llbracket e \rrbracket(\rho) = \mathbf{true} \Rightarrow \rho \in \gamma_n(\text{guard}(e, d))$.

We let $\mathcal{I}^\sharp \in D^\sharp$ be a sound approximation of the traces made of just one initial state, i.e. $\{\langle \sigma \rangle \mid \sigma \in \mathcal{I}\} \subseteq \gamma(\mathcal{I}^\sharp)$; we let $\mathcal{F}^\sharp \in D^\sharp$ be a sound approximation of \mathcal{F} in the same way, where \mathcal{F} may be either \mathcal{F}_l or $\mathcal{F}_{\mathfrak{D},l}$. The purpose of the following subsections is to approximate the semantic slice T .

2.3 Forward Analysis

We consider here the approximation of \overrightarrow{T} . It is well-known that a sound abstract interpreter in D^\sharp can be derived from \overrightarrow{F} . More precisely, we can define a family of functions $\overrightarrow{\delta}_{l,l'} : D_n^\sharp \rightarrow D_n^\sharp$ that compute the effect of any transition at the abstract level. The soundness of $\overrightarrow{\delta}_{l,l'}$ writes $\forall \rho, \rho' \in \mathbb{M}, \forall d \in D_n^\sharp, \rho \in \gamma_n(d) \wedge (l, \rho) \rightarrow (l', \rho') \Rightarrow \rho' \in \gamma_n(\overrightarrow{\delta}_{l,l'}(d))$ and is a direct consequence of the soundness of the basic abstract operations. The forward abstract interpreter is: $\overrightarrow{F}^\sharp : D^\sharp \rightarrow D^\sharp; I \mapsto \lambda(l \in \mathbb{L}). \sqcup \{ \overrightarrow{\delta}_{l,l'}(I(l)) \mid l' \in \mathbb{L} \}$ (this presentation leaves the choice for an iteration strategy; ASTRÉE uses a denotational iteration scheme, so as to keep the need for local invariant storage down). The soundness of the forward abstract interpreter is proved by standard abstract interpretation methods [8].

Theorem 1 (Soundness of the forward abstract interpreter). $T \subseteq \overrightarrow{T} \subseteq \gamma(\mathbb{l}_0)$ where $\mathbb{l}_0 = \mathbf{lfp}_{T^\sharp} \overrightarrow{F}^\sharp$.

Example 2 (Ex. 1 continued). A simple non relational analysis yields the invariant $b \in \{\mathbf{true}, \mathbf{false}\}$, $y \geq 0$ at point l_3 ; the assertion is not proved safe.

2.4 The Backward Interpreter

We consider the refinement of the approximation \mathbb{l}_0 of \overrightarrow{T} (hence, of T) into a better approximation, by taking into account the second fixpoint \overleftarrow{T} .

A straightforward way to do this would be to design a backward interpreter in the same way as we did for $\overrightarrow{F}^\sharp$ and to compute the intersection of both analyses. Yet, this approach would not be effective, mainly because in most cases, the greatest pre-conditions are not very precise, so that we would face a major loss of precision. For instance, in the case of a function call through a pointer dereference $(\star \mathbf{f})()$, the flow depends on the value of \mathbf{f} *before* the call; hence, the called function cannot be determined from the state after the call and the backward analysis of such a statement with no data about the pre-condition would be very imprecise ($\star \mathbf{f}$ could be any function in the program). Examples of similar issues when analyzing assignments are given in Sect. 2.5.

Hence, the refining backward interpreter $\overleftarrow{F}_r^\sharp$ takes two elements of D^\sharp as inputs: an invariant to refine and an invariant to propagate backwards. It is based on a family of backward transfer functions $\overleftarrow{\delta}_{l,l'} : D^\sharp \times D^\sharp \rightarrow D^\sharp$ maps a pre-condition to refine and a post-condition into a refined pre-condition, as stated by the soundness condition: $\forall \rho, \rho' \in \mathbb{M}, \forall d, d' \in D_n^\sharp, \rho \in \gamma_n(d) \wedge \rho' \in \gamma_n(d') \wedge (l, \rho) \rightarrow (l', \rho') \Rightarrow \rho \in \gamma_n(\overleftarrow{\delta}_{l,l'}(d, d'))$ (i.e. d is refined into a stronger pre-condition, by taking into account the post-condition d'). The definition for a very simple $\overleftarrow{\delta}_{l,l'}$ operator is given and discussed below. It is based on a backward abstract assignment operator $\overleftarrow{\text{assign}} : \mathbb{X} \times \mathbb{e} \times D_n^\sharp \times D_n^\sharp \rightarrow D_n^\sharp$, satisfying a similar soundness condition. The design of this operator is detailed in Sect. 2.5.

- assignment $l_0 : x := e$; $l_1 : \overleftarrow{\delta}_{l_0, l_1}(d_0, d_1) = \overleftarrow{\text{assign}}(x, e, d_0, d_1)$
- conditional $l_0 : \mathbf{if}(e)\{l_1^t : s^t; l_2^t\}\mathbf{else}\{l_1^f : s^f; l_2^f\}$; l_3 :
 $\overleftarrow{\delta}_{l_0, l_1^t}(d_0, d_1) = d_0 \sqcap d_1$ $\overleftarrow{\delta}_{l_2^t, l_3}(d_2, d_3) = d_2 \sqcap d_3$
 $\overleftarrow{\delta}_{l_0, l_1^f}(d_0, d_1) = d_0 \sqcap d_1$ $\overleftarrow{\delta}_{l_2^f, l_3}(d_2, d_3) = d_2 \sqcap d_3$
- loop $l_0 : \mathbf{while}(e)\{l_1 : s; l_2\}$; l_3 :
 $\overleftarrow{\delta}_{l_0, l_1}(d_0, d_1) = d_0 \sqcap d_1$ $\overleftarrow{\delta}_{l_2, l_0}(d_2, d_0) = d_2 \sqcap d_0$ $\overleftarrow{\delta}_{l_0, l_3}(d_0, d_3) = d_0 \sqcap d_3$
- assertion $l_0 : \mathbf{assert}(e)$; $l_1 : \overleftarrow{\delta}_{l_0, l_1}(d_0, d_1) = d_0 \sqcap d_1$
- input $l_0 : \mathbf{input}(x)$; $l_1 : \overleftarrow{\delta}_{l_0, l_1}(d_0, d_1) = d_0 \sqcap \text{forget}(x, d_1)$

It might be desirable to improve the precision by locally refining the computation of $\overleftarrow{\delta}_{l, l'}$. Indeed, if $\overleftarrow{\delta}_{l, l'}$ and $\overleftarrow{\delta}_{l', l''}$ are sound, then so is $\overleftarrow{\delta}_{l, l''}^{(n)} : (d, d') \mapsto d^{(n)}$, where: $d^{(0)} = \overleftarrow{\delta}_{l, l'}(d, d')$ and $\forall n \in \mathbb{N}$, $d^{(n+1)} = \overleftarrow{\delta}_{l, l'}(d^{(n)}, \overleftarrow{\delta}_{l', l''}(d^{(n)}))$. This process is known as local iterations [17] and usually allows to improve the precision of backward abstract operations and condition testings. For instance, in the case of the **if** statement, we may replace $\overleftarrow{\delta}_{l_0, l_1^t}$ with $\overleftarrow{\delta}_{l_0, l_1^t}(d_0, d_1) = \text{guard}(e, d_0 \sqcap d_1)$. Our experience proved local iterations not extremely useful, in the presence of a refined abstract domain, able to carry out rather expressive constraints.

The backward analyzer is defined by a function $\overleftarrow{F}_r^\# : D^\# \times D^\# \rightarrow D^\#$; $(I, I') \mapsto \lambda(l \in \mathbb{L}). \sqcup \{ \overleftarrow{\delta}_{l, l'}(I(l), I'(l')) \mid l' \in \mathbb{L} \}$ and satisfies the soundness result:

Theorem 2 (Soundness of the backward abstract interpreter). $T \subseteq \gamma(\mathbb{l}_1)$ where $\mathbb{l}_1 = \mathbf{lfp}_{\mathcal{F}^\# \sqcap \mathbb{l}_0}^\# [\lambda(I \in D^\#). \overleftarrow{F}_r^\#(\mathbb{l}_0, I)]$ and \mathbb{l}_0 is the result of the forward analysis (Th. 1).

2.5 Backward Assignment

The Domain. ASTRÉE uses a reduced product of several domains, including an interval domain, constraints among boolean variables or between boolean and scalar variables. Among the numerical relational domains, we can cite octagons [25] that express relations of the form $\pm x \pm y \leq c$ and specific domains like [13], adapted to the analysis of control command software components. Complex expressions can be abstracted prior to evaluation inside the abstract domain into interval linear forms [24]: given an abstract value $d \in D_n^\#$, e is abstracted into $e' = \mathbf{lin}(e, d) = \sum_k I_k \cdot x_k$ (\cdot is a product operator, $\forall k$, I_k is a real interval, x_k a variable), such that $\forall \rho \in \gamma_n(d)$, $\llbracket e \rrbracket(\rho) \in \llbracket e' \rrbracket(\rho)$. Linearization allows complex (e.g. non linear) expressions to be analyzed precisely inside relational domains.

We consider now the definition of $\overleftarrow{\text{assign}}(x, e, d_{\text{pre}}, d_{\text{post}})$. Note that we assume that the l-value x is resolved exactly; this is always the case in the subset of \mathbb{C} introduced in Sect. 2.1. In practice, may-assign (e.g. in the case of arrays) and assignment of pointer values are also taken into account. In the proofs below, we let $\rho \in \gamma_n(d_{\text{pre}})$; we write $v = \llbracket e \rrbracket(\rho)$ and we also assume $\rho[x \leftarrow v] \in \gamma_n(d_{\text{post}})$.

Boolean Transfer Function. If x is a boolean variable, we let:

$$\overleftarrow{\text{assign}}(x, e, d_{\text{pre}}, d_{\text{post}}) = \left\{ \begin{array}{l} \text{guard}(e, \text{forget}(x, \text{guard}(x, d_{\text{post}})) \sqcap d_{\text{pre}}) \\ \sqcup \text{guard}(-e, \text{forget}(x, \text{guard}(-x, d_{\text{post}})) \sqcap d_{\text{pre}}) \end{array} \right.$$

Indeed, let us assume $v = \mathbf{true}$. Then $\rho \in \gamma_n(\text{forget}(x, \text{guard}(x, d_{\text{post}})))$, due to the hypothesis on $\rho[x \leftarrow \mathbf{true}]$. Moreover, $\llbracket e \rrbracket(\rho) = \mathbf{true}$, so $\rho \in \gamma_n(\text{guard}(e, \text{forget}(x, \text{guard}(x, d_{\text{post}}))))$, which shows the soundness of the above definition.

Arithmetic Backward Transfer Function. Let us assume now that x has scalar type, e.g. floating point. We let $\mathbf{lin}(e, d) = \sum_k I_k \cdot x_k$ be an interval linear form of e . We consider the derivation of new octagon and interval constraints:

- the octagon domain provides backward assignment and guard abstract transfer functions for interval linear form expressions [25];
- the interval information in d_{pre} is refined as follows: we let $\mathfrak{I}_x^{\text{pre}}$ (resp. $\mathfrak{I}_x^{\text{post}}$) denote the interval information about x in d_{pre} (resp. d_{post}) and we compute a refined interval information $\mathfrak{I}_{x_j}^{\text{ref}}$ for x_j . The soundness of linearization implies that $v \in (\sum_{k \neq j} I_k \cdot \mathfrak{I}_{x_k}^{\text{pre}}) + I_j \cdot \rho(x_j)$; hence, if $0 \notin I_j$, $\rho(x_j) \in (v - (\sum_{k \neq j} I_k \cdot \mathfrak{I}_{x_k}^{\text{pre}})) / I_j \subseteq (\mathfrak{I}_x^{\text{post}} - (\sum_{k \neq j} I_k \cdot \mathfrak{I}_{x_k}^{\text{pre}})) / I_j$, so that the definition $\mathfrak{I}_{x_j}^{\text{ref}} = (\mathfrak{I}_x^{\text{post}} - (\sum_{k \neq j} I_k \cdot \mathfrak{I}_{x_k}^{\text{pre}})) / I_j$ is correct. These refined intervals are computed with a floating point-based approximation of the semantics of linear interval forms defined in terms of real numbers [24].

Note that the linearization should be computed using d_{pre} : using d_{post} would be unsound, since the value of the assigned variable changed; d_{pre} is also most useful to get the interval information *before* the assignment; hence, the first argument of $\overleftarrow{\text{assign}}$ is crucial to compute a precise and sound d_{pre} .

Example 3 (Backward assignment for intervals). We consider the assignment $x := y \cdot x + z$, $d_{\text{pre}} = \{x \geq 0, y \in [1, 2], z \in [1, 2], \dots\}$, $d_{\text{post}} = \{x \in [3, 4], \dots\}$. Linearization converts it into $x := [1, 2] \cdot x + z$; the backward assignment refines the range for x into $[0.5, 3]$. Local iteration would not improve the precision.

2.6 Iteration Strategies

Iterative Refining Process. At the concrete level, T could be defined as the intersection of two independent fixpoints. However, at the abstract level, the invariant $\mathbb{1}_1$ obtained after one forward analysis and one backward analysis might be refined by further analyses. For instance, in case the backward analysis reveals that no trace is going through the true branch of a conditional $l : \mathbf{if}(e)\{s_t\}\mathbf{else}\{s_f\}; l' : s'$, a refining forward analysis from $\mathbb{1}_1$ may refine the local invariants inside s' , since the possible imprecision due to the least upper bound at l' no longer occurs. Note that a further backward analysis would likely improve the results inside s_f also.

Therefore, we propose to define a refining forward analysis and to iterate the refining forward-backward process [6,9]. The *refining forward analyzer* $\overrightarrow{F}_r^\sharp$:

$D^\sharp \times D^\sharp \rightarrow D^\sharp$ is based on the forward analyzer and refines its first argument as the backward analyzer: $\vec{F}_r^\sharp : (I, I') \mapsto \lambda(l' \in \mathbb{L}). \sqcup \{ \vec{\delta}_{l,l'}(I(l)) \sqcap I(l') \mid l \in \mathbb{L} \}$.

The *refining sequence* $(\mathbb{l}_n)_{n \in \mathbb{N}}$ is defined by:

- \mathbb{l}_0 has been defined in Th. 1 by $\mathbb{l}_0 = \mathbf{lfp}_{\mathcal{I}^\sharp}^\sharp \vec{F}^\sharp$;
- if $n \geq 0$, $\mathbb{l}_{2n+1} = \mathbf{lfp}_{\mathcal{F}^\sharp \sqcap \mathbb{l}_{2n}}^\sharp [\lambda(I \in D^\sharp). \overleftarrow{F}_r^\sharp(\mathbb{l}_{2n}, I)]$ (akin to \mathbb{l}_1 , see Th. 2);
- if $n \geq 0$, $\mathbb{l}_{2n+2} = \mathbf{lfp}_{\mathcal{I}^\sharp \sqcap \mathbb{l}_{2n+1}}^\sharp [\lambda(I \in D^\sharp). \vec{F}_r^\sharp(\mathbb{l}_{2n+1}, I)]$.

Theorem 3 (Soundness of the forward-backward refinement). *The above process is sound: $\forall n \in \mathbb{N}, T \subseteq \gamma(\mathbb{l}_n)$.*

The proof is done by induction; it is similar to [9, Chap. 6]. Note that, given \mathcal{I}^\sharp and \mathcal{F}^\sharp , the sequence of refined invariants is obtained *fully automatically*.

Example 4 (Ex. 1 continued: refined analysis). We let \mathcal{F}^\sharp be $x \in [-10, 10] \wedge y \geq 0 \wedge b = \mathbf{true}$; clearly $\mathcal{F}_{\mathfrak{D}, l_3} \subseteq \gamma_n(\mathcal{F}^\sharp)$. The table below shows the result of the first two refining iterations, using a non relational abstraction:

label	\mathbb{l}_0	\mathbb{l}_1	\mathbb{l}_2
l_1	\top	\perp	\perp
l_2	$y \geq 0$	$-10 < x < 10 \wedge y \geq 10$	\perp
l_3	$y \geq 0 \wedge b \in \{\mathbf{true}, \mathbf{false}\}$	$-10 < x < 10 \wedge y \geq 10 \wedge b = \mathbf{true}$	\perp

$T_{\mathfrak{D}, l_3} = \emptyset$: the second refining iteration proves the correctness of the program, i.e. the alarm was false (Sect. 2.2, Case a).

Local Iterations. The above refinement process is not optimal from the efficiency point of view. In the case of the **if** statement considered above, it amounts to completing the backward analysis of the *whole* program before doing a new forward analysis so as to refine the invariant at label l . We might want to do *local iterations*, that is carrying out forward and backward local analysis steps in a single iteration phase. In practice, we found that the refinement process done with an expressive abstract domain (like the domain present in *ASTRÉE*) does not require much local iterations. Carrying out iterative refinements on large blocks of code (e.g. functions) was a more efficient strategy.

Implementation of the Interpreters. The forward analyzer *ASTRÉE* is written as a function that inputs a statement and an abstract pre-condition and returns an abstract post-condition; it is defined in denotational style, recursively on the syntax. The export of invariants is optional and one may choose the labels local invariants are exported at. The refining forward analyzer is based on the latter; a parameter just forces it to compute greater lower bounds with a previously computed invariant. The backward analyzer is very similar (same layout, same iteration strategy).

3 Specifying Alarm Contexts

Sect. 2 described the forward-backward analysis involved in the approximation of the set of “dangerous traces”. Yet, it does not solve the following issues:

- if we analyze backwards a statement $\mathbf{while}(e)\{l : \mathbf{assert}(e); \dots\}$, the backward interpreter computes a least-fixpoint on the loop; at the end of the process the invariant at l approximates not only the states right before an error occurs but also the states encountered one, two, or many iterations before, which results in a massive loss of precision at l ;
- after refinement of the invariants, we may have the intuition that $T_{\mathfrak{D},l} \neq \emptyset$; should that case arise, we would like to envisage and check an “error scenario”, which needs to be defined.

Example 5. We consider the example shown on Fig. 1(b) along this section. Clearly, this program may fail: it may read a negative input; at the next iteration, y is negative, which causes the assertion to fail. However, if the input is always positive, it does not fail. Last, note that it will not fail in the first iteration. The attempt to determine the alarm raised after computing \mathbb{I}_0 using a simple interval analysis leaves us in the undetermined case (Sect. 2.2, Case c).

3.1 Restriction to an Execution Pattern

We propose to extend the semantic slicing introduced in Sect. 2, by specifying some *execution patterns* in addition to the initial and final states: for instance, we may restrict a slice to the executions that cause an error after at least two iterations of the main loop and distinguish the states encountered in the last two iterations when approximating this slice. In practice, we resort to some kind of trace partitioning technique, that fits in the framework of [23].

Restriction to a Pattern. We extend the syntax of the language presented in Sect. 2.1 with a statement $l_0 : \mathbf{cnt}; l_1$. The new semantics should keep track of the order such statements are executed in. We propose to abstract this order.


Our approach involves the choice of an *automaton* $(\mathbb{Q}, \rightsquigarrow)$, where \mathbb{Q} is a finite set of states and $(\rightsquigarrow) \subseteq \mathbb{Q} \times \mathbb{L} \times \mathbb{Q}$ is a transition relation (we write $q \xrightarrow{l} q'$ for $(q, l, q') \in (\rightsquigarrow)$). The labels are replaced with pairs $(l, q) \in \mathbb{L} \times \mathbb{Q}$ in the definition of the concrete semantics: we replace \mathbb{L} with $\mathbb{L}_p = \mathbb{L} \times \mathbb{Q}$; \mathbb{S} with $\mathbb{S}_p = \mathbb{L}_p \times \mathbb{M} \cup \{\Omega\}$. The new, partitioned semantics $\llbracket P \rrbracket_p$ is defined similarly to $\llbracket P \rrbracket$, using the new transition relation $(\rightarrow_p) \subseteq \mathbb{S}_p \times \mathbb{S}_p$ instead of (\rightarrow) :

- case of $l_0 : \mathbf{cnt}; l_1$: if $q_0 \xrightarrow{l_0} q_1$, then $\forall \rho \in \mathbb{M}, ((l_0, q_0), \rho) \rightarrow_p ((l_1, q_1), \rho)$;
- case of any other transition:
 - if $(l_0, \rho_0) \rightarrow (l_1, \rho_1)$, then $((l_0, q), \rho_0) \rightarrow_p ((l_1, q), \rho_1)$;
 - if $(l_0, \rho_0) \rightarrow \Omega$, then $((l_0, q), \rho_0) \rightarrow_p \Omega$.

The *execution pattern* defined by a pair of states (q, q') is $\gamma_{\mathbb{Q}}(q, q') = \{((l_0, q_0), \rho_0), \dots, ((l_n, q_n), \rho_n) \mid q_0 = q \wedge q_n = q'\}$.

Example 6 (Ex. 5 continued). We insert a \mathbf{cnt} statement in the loop and consider the automaton \mathcal{Q} below. Then, $\gamma_{\mathbb{Q}}(q_0, q_2)$ specifies all the traces reaching the dangerous label at the iteration n where $n \geq 2$ and distinguishes the last two iterations (q_1 stands for iteration $n-1$; q_2 stands for iteration n). The automaton

allows us to restrict to the executions that spend more than one iteration in the loop (hence, that may cause an error).



We write $\pi_{\mathbb{L}} : \mathbb{L}_p \rightarrow \mathbb{L}$ (resp. $\pi_{\mathbb{S}} : \mathbb{S}_p \rightarrow \mathbb{M}, \pi : \mathbb{S}_p^* \rightarrow \mathbb{S}^*$) for the erasure function that removes the elements of \mathbb{Q} in labels (resp. stores, traces); we let π also be defined for sets of traces. If $\tau \in \llbracket P \rrbracket_p$, then $\pi(\tau) \in \llbracket P \rrbracket$ (proof obvious).

Refining the Semantic Slice. We also need to extend \mathcal{I} and \mathcal{F} . Let $q_i, q_f \in \mathbb{Q}$. We define $\mathcal{I}_p = \{((l, q_i), \rho) \mid (l, \rho) \in \mathcal{I}\}$ and $\mathcal{F}_p = \{((l, q_f), \rho) \mid (l, \rho) \in \mathcal{F}\}$. The automaton $(\mathbb{Q}, \rightsquigarrow)$ and the states q_i, q_f are currently chosen by the user so as to specify some set of execution paths and to specialize even more the semantic slice T ; the automatic selection of refinements is left as future work (see discussion in Sect. 5). Other choices for \mathcal{I} or \mathcal{F} , involving several states in the automaton are possible (the extension is easy). The definition of T_p from $\mathcal{I}_p, \mathcal{F}_p$ is similar to the definition of T from \mathcal{I}, \mathcal{F} (Sect. 2.2). It satisfies the following property, which clearly shows that it is restricted to the execution pattern defined by q_i, q_f :

$$\pi(T_p) = T \cap \pi(\gamma_{\mathbb{Q}}(q_i, q_f))$$

Approximation of the Semantic Slice. We replace $D^\# = \mathbb{L} \rightarrow D_n^\#$ with the partitioning domain $D_p^\# = \mathbb{L} \times \mathbb{Q} \rightarrow D_n^\#$; we let $\gamma_p : D_p^\# \rightarrow \mathcal{P}(\mathbb{S}_p), I \mapsto \{((l_0, q_0), \rho_0), \dots, ((l_n, q_n), \rho_n) \mid \forall i, \rho_i \in \gamma_n(I(l_i, q_i))\}$. The definition of forward and backward abstract interpreters and of the sequence of refined invariants \mathbb{I}_n^p follows the steps of Sect. 2, with extended definitions for $\overrightarrow{\delta}_{l, l'}$, $\overleftarrow{\delta}_{l, l'}$:

- Case of $l_0 : \mathbf{cnt}; l_1$:
 - forward analysis: if $q_0 \rightsquigarrow^{l_0} q_1$, $\overrightarrow{\delta}_{(l_0, q_0), (l_1, q_1)}(d) = d$;
 - backward analysis: if $q_0 \rightsquigarrow^{l_0} q_1$, $\overleftarrow{\delta}_{(l_0, q_0), (l_1, q_1)}(d, d') = d \sqcap d'$;
- Case of other statements: $\overrightarrow{\delta}_{(l_0, q), (l_1, q)}$ (resp. $\overleftarrow{\delta}_{(l_0, q), (l_1, q)}$) is defined like $\overrightarrow{\delta}_{l_0, l_1}$ (resp. $\overleftarrow{\delta}_{l_0, l_1}$) in Sect. 2.

Theorem 4 (Soundness). *The static analysis of the partitioned system leads to a sequence of sound invariants: $\forall n \in \mathbb{N}, T_p \subseteq \gamma_p(\mathbb{I}_n^p)$.*

Proof: with respect to $\llbracket P \rrbracket_p$.

Example 7 (Execution patterns). Many patterns can be encoded easily (we assume that the program is of the form $\mathbf{while}(e)\{l_c : \mathbf{cnt}; \dots\}$):




Fig. 2. Automata specifying trace patterns

- On Fig. 2(a), q_e and q_o correspond to even and odd iteration numbers: $\gamma_{\mathcal{Q}}(q_e, q_o)$ slices the traces iterating the loop an even number of times; it also helps distinguishing states reached after an odd (resp. even) number of iterations;
- On Fig. 2(b), $\gamma_{\mathcal{Q}}(q_0, q_1)$ corresponds to the first iteration, $\gamma_{\mathcal{Q}}(q_0, q_2)$ to the n -th iteration ($n \geq 2$).

Example 8 (Ex. 5 continued).

- The refinement of $T_{\mathcal{D},(l_d, q_1)}$ with the automaton of Fig. 2(b) and the pattern $q_i = q_0, q_f = q_1$ shows that no error may happen in the first iteration;
- Similarly, the refinement of $T_{\mathcal{D},(l_d, q_2)}$ with the automaton \mathcal{Q} (Ex. 6) gives some intuition about the traces that cause an error: at (l_d, q_2) , we get $y \leq 0$; at (l_d, q_1) , we get $x \leq 0$, which suggests the input of x should be negative one iteration before the error. We wish now to verify this error scenario.

Remarks. Note that the choice of an automaton with only one state q and such that for any statement $l : \mathbf{cnt}, q \xrightarrow{l} q$ results in the same analyses as in Sect. 2.

The trace partitioning presented in this section runs on the top of the one described in [23]; the latter aims at computing more precise invariants thanks to delayed merges of flows (e.g. out of **while** or **if** statements). Our goal here is to extract some execution patterns and to refine the corresponding invariants.

3.2 Restriction to a Set of Inputs

We now consider the slices defined by constraining the inputs; for instance, this may allow to show that this input always leads to an error.

Specification of Inputs. We let \mathbb{L}_{in} denote the set of **input** statements labels: $\mathbb{L}_{\text{in}} = \{l \in \mathbb{L} \mid l : \mathbf{input}(x_l)\}$. An *input specification* is a function $\nu : \mathbb{L}_p \rightarrow \mathcal{P}(\mathbb{V})$, mapping a label to the set of values that may be read at this point. The definition of ν over the partitioned labels allows to select different inputs for different execution contexts (corresponding to different states in the automaton introduced in Sect. 3.1) at the same label. The denotation of the input function ν is the set of traces $\gamma_{\mathbb{V}}(\nu) = \{((l_0, q_0), \rho_0), \dots, ((l_n, q_n), \rho_n) \mid \forall i, l_i \in \mathbb{L}_{\text{in}} \Rightarrow \rho_{i+1}(x_{l_i}) \in \nu(l_i, q_i)\}$: such traces satisfy the property that reading an input at label (l, q) yields a value in $\nu(l, q)$.

Refining the Semantic Slice. The semantic slice constrained to ν is:

$$T_v = T_p \cap \gamma_{\mathbb{V}}(\nu) = T \cap \gamma_{\mathbb{Q}}(q_i, q_f) \cap \gamma_{\mathbb{V}}(\nu) .$$

Approximation of the Semantic Slice. The only modification required to take into account the input specification concerns the rule for the $l_0 : \mathbf{input}(x); l_1$ statement. In this case, we let $\overrightarrow{\delta}_{(l_0, q), (l_1, q)}(d) = \mathit{forget}(x, d) \sqcap \nu^\sharp(l_0, q)$ where $\nu^\sharp(l_0, q)$ is a sound approximation of $\nu(l_0, q)$: $\{\rho \in \mathbb{M} \mid \rho(x) \in \nu(l_0, q)\} \subseteq \gamma_n(\nu^\sharp(l_0))$. The case of the backward analysis requires no modification.

Theorem 5 (Soundness). *The resulting abstract interpreters are sound and compute a sequence of sound refined invariants: $\forall n \in \mathbb{N}, T_v \subseteq \gamma_p(\mathbb{I}_n^p)$.*

The proof follows the definition of a variation $\llbracket P \rrbracket_v$ of the concrete semantics $\llbracket P \rrbracket_p$: $\llbracket P \rrbracket_v$ is obtained from $\rightarrow_{p, \nu}$ just as $\llbracket P \rrbracket_p$ from \rightarrow_p , where $\rightarrow_{p, \nu}$ is the transition relation constrained by the input function ν . The only modification in the definition of $\rightarrow_{p, \nu}$ comes from the case of the $l : \mathbf{input}(x_l); l'$ statement: $((l, q), \rho) \rightarrow ((l', q), \rho[x \leftarrow v])$ where $v \in \nu(l, q)$.

Example 9 (Ex. 5 continued). Let us consider the input specification $\nu(l_i, q_1) = -1$. Then, \mathbb{I}_0 shows that $y = -1$ at point (l_d, q_2) (the interval analysis proves this property). Hence, the automaton \mathcal{Q} and the input specification ν define a valid error scenario: any execution iterating the loop n times and such that the value read for x during the $(n - 1)$ th iteration is -1 will result in an error. Such situations are feasible, so the static analysis showed a real bug in the program (Sect. 2.2, Case b).

Example 10 (Ex. 7 continued). The automaton of Fig. 2(a) allows to specify a cyclic input; the automaton of Fig. 2(b) allows to isolate initialization inputs read during the first iteration and inputs read at iteration n for $n \geq 2$.

Currently, the function ν^\sharp should be provided by the user; further work should allow to synthesize an input specification ν^\sharp exhibiting an error.

4 Slicing

The size of the program is a major issue when computing semantic slices as suggested in Sect. 2.4 and 2.6. Indeed, forward-backward analyses require saving local invariants, which would induce a dramatic memory cost, if applied to the whole program. Therefore, we propose to use regular, syntactic slicing techniques [29,19] so as to restrict the amount of code to apply the refining analyses to. We use the notations of Sect. 2 for the sake of simplicity (even though mixing slicing techniques and the methods introduced in Sect. 3 is straightforward).

We assume a program s is given, that contains a statement $l_d : \mathbf{assert}(e)$. We write $\mathbf{use}(e)$ for the set of variables that appear in expression e .

Slicing. A slicing criterion is a set $\mathcal{C} \subseteq \mathbb{L} \times \mathbb{X}$ of pairs made of a label and a variable ; it specifies a set of variables we wish to observe the value of, at some point. A typical choice is $\mathcal{C} = \{(l_d, x) \mid x \in \mathbf{use}(e)\}$. See eg. [19] for standard slicing algorithms. Slicing is sound in the sense that it does not remove any behavior of the original program for some observation including the dependence closure of the criterion. As a consequence, the static analysis of the slice yields a safe approximation of the semantic slice of the initial program. Beyond the restriction of the size of the code to analyze, an advantage of considering the slice defined by the $l_d : \mathbf{assert}(e)$ statement is that most of the remaining statements and variables are relevant to the observation at label l_d , i.e. to the alarm under investigation.

Reducing the Size of Slices. If $\mathbb{I}_n(l) = \perp$, then \mathbb{I}_n proves that the statement at label l is not relevant to the semantic slice of interest T ; hence, this statement can be safely removed from the slice and its dependences thrown away, which allows to reduce even more the size of the syntactic slice. Such a transformation preserves the soundness and should speed up the computation of \mathbb{I}_{n+k} ($k \geq 1$).

Approximation of Slices. Slicing should reduce the programs to analyze to a reasonable size; however, even the slices extracted from some **assert** statements may have prohibitive sizes, when extracted from very large programs, with long, cyclic dependence chains. Thus, we propose to do “aggressive slicing” and to approximate any removed statement in a sound manner during the analysis.

For instance, let us consider the forward analysis of a statement $l_0 : x := e; l_1$ (that should be extracted in the slice). As seen in Sect. 2, the forward abstract transfer function for this statement is $\vec{\delta}_{l_0, l_1} : d \mapsto \overrightarrow{\mathbf{assign}}(x, e, d)$. The aggressive slicing of this statement consists in replacing the previous definition of $\vec{\delta}_{l_0, l_1}$ with the following: $\vec{\delta}_{l_0, l_1} : d \mapsto \overrightarrow{\mathbf{forget}}(x, d)$. Observe that this is sound: $d \mapsto \overrightarrow{\mathbf{forget}}(x, d)$ approximates all the concrete transitions defined by the assignment; hence, this new definition for $\vec{\delta}_{l_0, l_1}$ leads to a sound forward and backward abstract interpreters (the soundness results of Th. 1, Th. 2, and Th. 3 are preserved). Among the possible strategies to reduce the size of “aggressive slices”, we can cite the limiting of dependency chains, the restriction to a given subset of variables or the elimination of loop-carried dependences: these approaches lead to an under-approximation $\widehat{\mathcal{C}}$ of the dependences induced by \mathcal{C} .

5 Case Studies

A typical alarm investigation session proceeds as follows:

1. do a forward analysis, determine a superset of the possible errors (Th. 1);
2. choose an alarm to investigate; restrict to a slice including the alarm point;
3. define $\mathcal{I}^\sharp, \mathcal{F}^\sharp$, attempt to prove the alarm wrong with forward-backward refinement (Th. 3), otherwise a more precise alarm context slice is found;
4. in case of failure, specialize the alarm context (Sect. 3.1);
5. in case no attempt to get the analyzer to prove $T_{\mathcal{D}, l} = \emptyset$ succeeds, then attempt to prove the alarm true by choosing a set of inputs (Sect. 3.2).

Application to a Family of Programs. We applied our methodology to the alarms raised by ASTRÉE on a series of 3 early development versions of some critical embedded programs (bugs were not unlikely in the development versions).

Size of the C code (lines)	67 500	233 000	412 000
Number of functions	650	1900	2900
Analysis time (\mathbb{l}_0) in sec.	1 300	16 200	37 500
Number of alarms	4	1	0
Alarm names	a_1, a_2, a_3, a_4	a_5	-

Slicing (Sect. 4) showed that a_2 (resp. a_4) is a direct consequence of a_1 (resp. a_3); hence, we restricted to the investigation of a_1 , a_3 and a_5 . The computation of a semantic slice for the corresponding dangerous states on the slices revealed rather informative conditions on the inputs. Specializing some inputs and carrying out a new, forward analysis *allowed to prove the alarms true*, thanks to an input specification as in Ex. 10. The table below provides some data about the process: the number of input constraints is the number of points an input constraint had to be specified for (Sect. 3.2); the number of execution patterns corresponds to the number of automata we considered (Sect. 3.1). The size of the slices (number of lines, functions and variables) involved in the alarms show that a_1, a_3 were rather subtle; a_5 was much simpler. The number of additional constraints generated during the forward-backward refinement is rather difficult to express simply due to the trace partitioning, and to the use of sophisticated numerical domains; we can only mention that it is much higher than the number of variables or of program points. One forward-backward iteration necessitates a reasonable amount of resources for these slices (up to 1 min., 80 Mb).

Alarm	a_1	a_3	a_5
Size of the slice (lines)	1280	4096	244
Number of functions in the slice	29	115	8
Number of variables in the slice including: int, bool, float variables	215	883	30
Execution patterns	15, 60, 146	122, 553, 208	7, 11, 23
Input constraints	2	2	2
Input constraints	4	4	2

The only manual step is the choice of adequate execution patterns and of constraints on inputs, so as to get an error scenario; in all the above cases, these numbers are very low, which shows the amount of work for the user is very reasonable: only 4 inputs had to be chosen in the most complicated case (a_3). However each of these choices had to be made carefully, with respect to complex conditions on bit-fields and arithmetic values. The choices for the execution patterns to examine only required considering very few simple automata (similar to unrolling of loops, akin to Fig. 2(b) and Ex. 6), so that the selection of execution patterns should be easy to automatize.

All alarms found involve intricate floating point computations. For instance, a_5 is due to a mis-use of (interpolated) trigonometric functions, leading to a possibly negative result, causing a square root computation to fail.

Early Experience Conclusions. The use of the system reduced the alarm investigation time to a few hours in the worst case we faced; the refining analyses are fully automatic and default parameters (fixed number of global forward-backward steps, no local iterations) did not have to be twicked too much to give good results. Fully manual inspection of such alarms would have required days of work and would have made the definition of an error scenario much more involved. Moreover, we could successfully classify all alarms, which means that *no false alarm remains*.

6 Conclusion and Related Work

We proposed a framework adapted to alarm inspection. Early experiments are positive about the ability of the system to reduce the burden of tracking the source of alarms in ASTRÉE: overall, *all considered alarms could be classified* (no case similar to Fig. 1(c) left), which is a very positive result.

Some forms of conditioned slicing [21,4] attack a similar problem. However, these methods are essentially based on a purely syntactic process, not only for the extraction but also for the shape of the result (a slice is a subset of the program statements [29]). Slicing has been employed for debugging tasks. Recent advances in this area led to the implementation of conditioned slicing tools like [14], that could be applied to testing and software debugging [18]. However, our system is able to produce *semantic* slices, i.e. to provide global information about a set of executions instead of a mere syntactic subset of the program; this is a major advantage when investigating complex errors. The downside is the use of more sophisticated algorithms; however, syntactic slicing alone would not help significantly the alarm inspection process in ASTRÉE.

The search for counter-examples and automatic refinement has long been a motivation in the model-checking-based systems, such as [5,2,26,27,16]. In particular, the automatic refinement process plays a great role in the determination of the set of predicates (i.e. abstract domain) needed for a precise analysis [1]. Our goal is to bring such methods in static analyzers like ASTRÉE for a different purpose, i.e. to solve the few, subtle alarms, after an already very precise analysis [3] (the construction of the domain requires no internal refinement process).

Forward-backward analysis schemes have been applied, e.g. in [20], to the inference of safety properties. Some static analysis systems have been extended with counter-examples search facilities: [15] relies on random test generation; [12] uses a symbolic under-approximation of erroneous traces and theorem proving. The main difference is that we chose to start with an over-approximation of erroneous traces until conditions on inputs are precise enough so that a counter-example could be found since the search space for counter-examples was huge in our case, due to the size of the programs. For instance, the systematic exploration

of paths as in [12] over length above 1000, with hundreds of variables would not work. Moreover, we allow abstract error scenario to be tested unlike [15,12]: this reduces the amount of input constraints to fix to a minimum. On the other hand, we leave the automatic generation of counter-examples as a future work.

Future work should make the process more automatic for attempting to discover an error scenario, by proposing input sequences and restricting to adapted alarm contexts (which are user provided in Sect. 3.1 and Sect. 3.2). We also plan to make the choice of slices to analyze (Sect. 4) more sensible, by using the result of the initial forward analysis, to choose which part of the invariant to refine.

Acknowledgments. We wish to thank deeply B. Blanchet, P. Cousot, R. Cousot, J. Feret, C. Hymans, L. Mauborgne, A. Miné, and D. Monniaux for comments on early version of this paper and discussions. We are also grateful to M. Sagiv for interesting discussions about related work.

References

1. T. Ball, R. Majumdar, T. Millstein, and S. Rajamani. Automatic predicate abstraction of C programs. In *PLDI*, 2001.
2. T. Ball, M. Naik, and S. Rajamani. From symptom to cause: Localizing errors in counterexample traces. In *POPL*, 2003.
3. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. A Static Analyzer for Large Safety Critical Software. In *PLDI*, 2003.
4. G. Canfora, A. Cimitille, and A. D. Lucia. Condition program slicing. *Information and Software Technology; Special issue on Program Slicing*, 1998.
5. E. M. Clarke, O. Grumberg, S. Jha, Y. Lu, and H. Veith. Counterexample-guided abstraction refinement. In *CAV*, 2000.
6. P. Cousot. *Méthodes itératives de construction et d'approximation de points fixes d'opérateurs monotones sur un treillis, analyse sémantique des programmes*. PhD thesis, 1978.
7. P. Cousot. Semantic foundations of program analysis. In S. Muchnick and N. Jones, editors, *Program Flow Analysis: Theory and Applications*, pages 303–342. Prentice-Hall, Inc., Englewood Cliffs, New Jersey, 1981.
8. P. Cousot and R. Cousot. Abstract Interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, 1977.
9. P. Cousot and R. Cousot. Abstract Interpretation and Application to Logic Programs. *Journal of Logic Programming*, 13(2–3):103–179, 1992.
10. P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. The ASTRÉE analyzer. In *ESOP*, 2005.
11. N. Dor, M. Rodeh, and M. Sagiv. CSSV: Towards a realistic tool for statically detecting all buffer overflows in C. In *PLDI*, 2003.
12. G. Erez. Generating counter examples for sound abstract interpretation. Master's thesis, 2004.
13. J. Feret. Static analysis of digital filters. In *ESOP*, 2004.
14. C. Fox, S. Danicic, M. Harman, and R. Hierons. ConSIT: A Conditioned Program Slicing System. *Software - Practice and Experience*, 2004.
15. F. Gaucher, E. Jahier, B. Jeannet, and F. Maraninchi. Automatic state reaching for debugging reactive programs. In *AADEBUG*, 2003.

16. R. Giacobazzi, F. Ranzato, and F. Scozzari. Making abstract interpretations complete. *Journal of the ACM*, pages 361–416, 2000.
17. P. Granger. Improving the results of static analyses programs by local decreasing iteration. In *FSTTCS*, 1992.
18. R. Hierons, M. Harman, C. Fox, , L. Ouarbya, and D. Daoudi. Conditioned slicing supports partition testing. *Journal of Software Testing, Verification and Reliability*, 2002.
19. S. Horwitz, T. Reps, and D. Binkley. Interprocedural Slicing using Program Dependence Graphs. *Programming Languages and Systems*, 1990.
20. B. Jeannet. Dynamic partitioning in linear relation analysis. *Formal Methods in System Design*, 2003.
21. B. Korel and J. Laski. Dynamic Program Slicing. *Information Processing Letters*, 1988.
22. T. Lev-Ami and M. Sagiv. TVLA: A system for implementing static analyses. In *SAS*, 2000.
23. L. Mauborgne and X. Rival. Trace Partitioning in Abstract Interpretation Based Static Analyzers. In *ESOP*, 2005.
24. A. Miné. Relational abstract domains for the detection of floating-point run-time errors. In *ESOP*, 2004.
25. A. Miné. *Weakly relational numerical abstract domains*. PhD thesis, 2004.
26. G. Pace, N. Halbwachs, and P. Raymond. Counter-example generation in symbolic abstract model-checking. In *6th International Workshop on Formal Methods for Industrial Critical Systems, FMICS*, 2001.
27. A. Podelski. Software model checking with abstraction refinement. In *VMCAI*, 2003.
28. A. Venet and G. Brat. Precise and efficient array bound checking for large embedded c programs. In *PLDI*, 2004.
29. M. Weiser. Program slicing. In *Proceeding of the Fifth International Conference on Software Engineering*, pages 439–449, 1981.