Abstraction-Guided Synthesis of Synchronization

Martin Vechev
IBM Research

Eran Yahav
IBM Research

Greta Yorsh
IBM Research

Abstract

We present a novel framework for automatic inference of efficient synchronization in concurrent programs, a task known to be difficult and error-prone when done manually.

Our framework is based on abstract interpretation and can infer synchronization for infinite state programs. Given a program, a specification, and an abstraction, we infer synchronization that avoids all (abstract) interleavings that may violate the specification, but permits as many valid interleavings as possible.

Combined with abstraction refinement, our framework can be viewed as a new approach for verification where both the program and the abstraction can be modified on-the-fly during the verification process. The ability to modify the program, and not only the abstraction, allows us to remove program interleavings not only when they are known to be invalid, but also when they cannot be verified using the given abstraction.

We implemented a prototype of our approach using numerical abstractions and applied it to verify several interesting programs.

Categories and Subject Descriptors D.1.3 [*Concurrent Programming*]; D.2.4 [*Program Verification*]
General Terms Algorithms, Verification
Keywords concurrency, synthesis, abstract interpretation

1. Introduction

We present *abstraction-guided synthesis*, a novel approach for synthesizing efficient synchronization in concurrent programs. Our approach turns the one dimensional problem of verification under abstraction, in which only the abstraction can be modified (typically via abstraction refinement), into a two-dimensional problem, in which *both the program and the abstraction can be modified* until the abstraction is precise enough to verify the program.

Based on abstract interpretation [10], our technique synthesizes a symbolic characterization of *safe schedules* for concurrent infinite-state programs. Safe schedules can be realized by modifying the program or the scheduler:

- Concurrent programming: by automatically inferring minimal atomic sections that prevent unsafe schedules, we assist the programmer in building correct and efficient concurrent software, a task known to be difficult and error-prone.
- Benevolent runtime: a scheduler that always keeps the program execution on a safe schedule makes the runtime system more reliable and adaptive to ever-changing environment and safety requirements, without the need to modify the program.

POPL'10, January 17–23, 2009, Madrid, Spain.

Given a program P, a specification S, and an abstraction function α, verification determines whether $P \models_\alpha S$, that is, whether P satisfies the specification S under the abstraction α. When the answer to this question is negative, it may be the case that the program violates the specification, or that the abstraction α is not precise enough to show that the program satisfies it.

When $P \not\models_\alpha S$, abstraction refinement approaches (e.g., [8, 3]) share the common goal of trying to find a finer abstraction α' such that $P \models_{\alpha'} S$. In this paper, we investigate a complementary approach, of finding a program P' such that $P' \models_\alpha S$ under the original abstraction α and P' admits a subset of the behaviors of P. Furthermore, we combine the two directions — refining the abstraction, and restricting program behaviors, to yield a novel abstraction-guided synthesis algorithm.

One of the main challenges in our approach is to devise an algorithm for obtaining such P' from the initial program P. In this paper, we focus on *concurrent programs*, and consider changes to P that correspond to restricting interleavings by adding synchronization.

Although it is possible to apply our techniques to other settings, concurrent programs are a natural fit. Concurrent programs are often correct on most interleavings and only miss synchronization in a few corner cases, which can be then avoided by synthesizing additional synchronization. Furthermore, in many cases, constraining the permitted interleavings reduces the set of reachable (abstract) states, possibly enabling verification via a coarser abstraction and avoiding state-space explosion.

The AGS algorithm, presented in Section 4, iteratively eliminates invalid interleavings until the abstraction is precise enough to verify the program. Some of the (abstract) invalid interleavings it observes may correspond to concrete invalid interleavings, while others may be artifacts of the abstraction. Whenever the algorithm observes an (abstract) invalid interleaving, the algorithm tries to eliminate it by either (i) modifying the program, or (ii) refining the abstraction.

To refine the abstraction, the algorithm can use any standard technique (e.g.,[8, 3]). These include moving through a predetermined series of domains with increasing precision (and typically increasing cost), or refining within the same abstract domain by changing its parameters (e.g., [4]).

To modify the program, we provide a novel algorithm that generates and solves *atomicity constraints*. Atomicity constraints define which statements have to be executed atomically, without an intermediate context switch, to eliminate the invalid interleavings. This corresponds to limiting the non-deterministic choices available to the scheduler. A solution of the atomicity constraints can be implemented by adding atomic sections to the program.

Our approach separates the process of identifying the space of solutions (generating the atomicity constraints) from the process of choosing between the possible solutions, which can be based on a quantitative criterion. As we discuss in Section 6, our approach provides a solution to a *quantitative synthesis* problem [5], as it

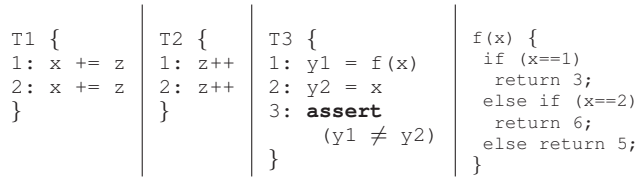

Figure 1. Simple example computing values of $y1$ and $y2$.

can compute a *minimally atomic* safe schedule for a program, a schedule that poses minimal atomicity constraints on interleavings, and does not restrict interleavings unnecessarily.

Furthermore, our approach can be instantiated with different methods for: (i) modifying the program to eliminate invalid interleavings (ii) refining the abstraction (iii) choosing optimal solutions (quantitative criterion) (iv) implementing the resulting solution.

The problem we address in this paper is closely related to the ones addressed by program repair [14, 12] and controller synthesis [20]. However, in contrast to these, our approach focuses on concurrent programs, uses abstract interpretation, and is able to handle infinite-state programs.

1.1 Main Contributions

The contributions of this paper can be summarized as follows:

- We provide a novel algorithm for inferring correct and efficient synchronization in concurrent programs. The algorithm infers minimal atomic sections that can be verified under a given abstraction.
- We advocate a new approach to verification where both the program and the abstraction can be modified on the fly during the verification process. This enables verification of a restricted program where verification of the original program fails.
- We implemented our approach in a prototype tool called GUARDIAN and applied it to synthesize synchronization for several interesting programs using numerical abstractions.

1.2 Limitations

Our focus in this paper is on the AGS algorithm (Sec. 4) and on an algorithm for eliminating invalid interleaving by adding atomic sections. While our approach can be instantiated with various abstraction-refinement algorithms and abstract domains, our current realization is quite modest:

- The abstraction-refinement approach we use in the paper is rather simplistic. Using more sophisticated refinement approaches is a topic of future work.
- We only implement a number of simple numerical abstract domains, which enable us to handle infinite-state numerical programs. To make the approach more widely applicable, we intend to integrate additional abstract domains in the future.

2. Overview

In this section, we demonstrate our technique on a simple illustrative example. The discussion in this section is mostly informal, additional formal details are provided in Section 4. Additional examples, inspired by real applications, are described in Section 7.

2.1 Example Program

Consider the example shown in Fig. 1. In this example, the program executes three processes in parallel: `T1||T2||T3`. Different interleavings of the statements executed by these processes lead to different values being assigned to $y1$ and $y2$. In every execution of the program there is a single value assigned to $y1$ and a single value assigned to $y2$. The assertion in `T3` requires that the values of $y1$ and $y2$ are not equal. Initially, the value of all variables are 0.

For example, $y1$ gets the value 6, and $y2$ gets the value 2 in the interleaving `z++;x+=z;x+=z;y1=f(x);y2=x;z++;assert`. In the interleaving `x+=z;x+=z;y1=f(x);y2=x;z++;z++;assert`, $y1$ gets the value 5, and $y2$ gets the value 0.

Fig. 2 (I) shows the possible values of $y1$ and $y2$ that can arise during *all possible* program executions, assuming that the macro f executes atomically. Note that in some interleavings $y1$ and $y2$ may be evaluated for different values of x (i.e., x can be incremented between the assignment to $y1$ and the assignment to $y2$). The point $y1 = y2 = 3$ (marked in red in Fig. 2 (I)) corresponds to values that violate the assertion. These values arise in the following interleaving: `z++; x+=z; y1=f(x); z++; x+=z; y2=x;assert`.

Our goal is to add efficient synchronization to the program such that the assertion in `T3` is not violated in any execution.

The AGS algorithm iteratively eliminates invalid interleavings (under an abstraction) by either modifying the program or the abstraction. Fig. 2 shows how the algorithm operates on the program of Fig. 1, and how it can move on both dimensions, choosing to modify either the program, or the abstraction, on every step. Before we explain Fig. 2, we explain how the algorithm modifies the program to eliminate invalid iterleavings without any abstraction.

2.2 Inferring Synchronization under Full Information

We begin by considering the example program without abstraction. Since this is an illustrative finite-state program, we can focus on the aspects of the algorithm related to generating atomicity constraints.

The algorithm accumulates atomicity constraints by iteratively eliminating invalid interleavings. Every invalid interleaving yields an atomicity constraint that describes *all possible ways* to eliminate that interleaving, by disabling context-switches that appear in it.

Under full information, the program of Fig. 1 has a single invalid interleaving `z++; x+=z; y1=f(x); z++; x+=z; y2=x; assert`. This interleaving can be eliminated by disabling either of the context switches that appear in this interleaving: the context switch between `x+=z` and `x+=z` in `T1`, between `z++` and `z++` in `T2`, and between `y1=f(x)` and `y2=x` in `T3`. This corresponds to the following atomicity constraint, generated by AGS algorithm:

$$\texttt{[y1=f(x),y2=x]} \vee \texttt{[x+=z,x+=z]} \vee \texttt{[z++,z++]}$$

This constraint is a disjunction of three atomicity predicates, of the form `[s1,s2]`, where `s1` and `s2` are consecutive statements in the program. Each atomicity predicate represents a context-switch that can eliminate the invalid interleaving, and the disjunction represents the fact that we can choose either one of these three to eliminate the invalid interleaving. For this program, there are no additional constraints, and any satisfying assignment to this constraint yields a correct program. For example, adding an atomic section around `z++` and `z++` in `T2` yields a correct program.

Since we can obtain multiple solutions, it is natural to define a quantitative criterion for choosing among them. This criterion can be based on the number of atomic sections, their length, etc. Our approach separates the process of identifying the space of solutions (generating the atomicity constraints) from the process of choosing between the possible solutions, which can be based on a quantitative criterion. In this example, each of the three possible solutions only requires a single atomic section of two statements.

Next, we illustrate how AGS operates under abstraction. In this example, we use simple numerical domains: parity, intervals, and octagon. In Section 7, we show refinement by increasing the set of variables for which the abstraction tracks correlations.

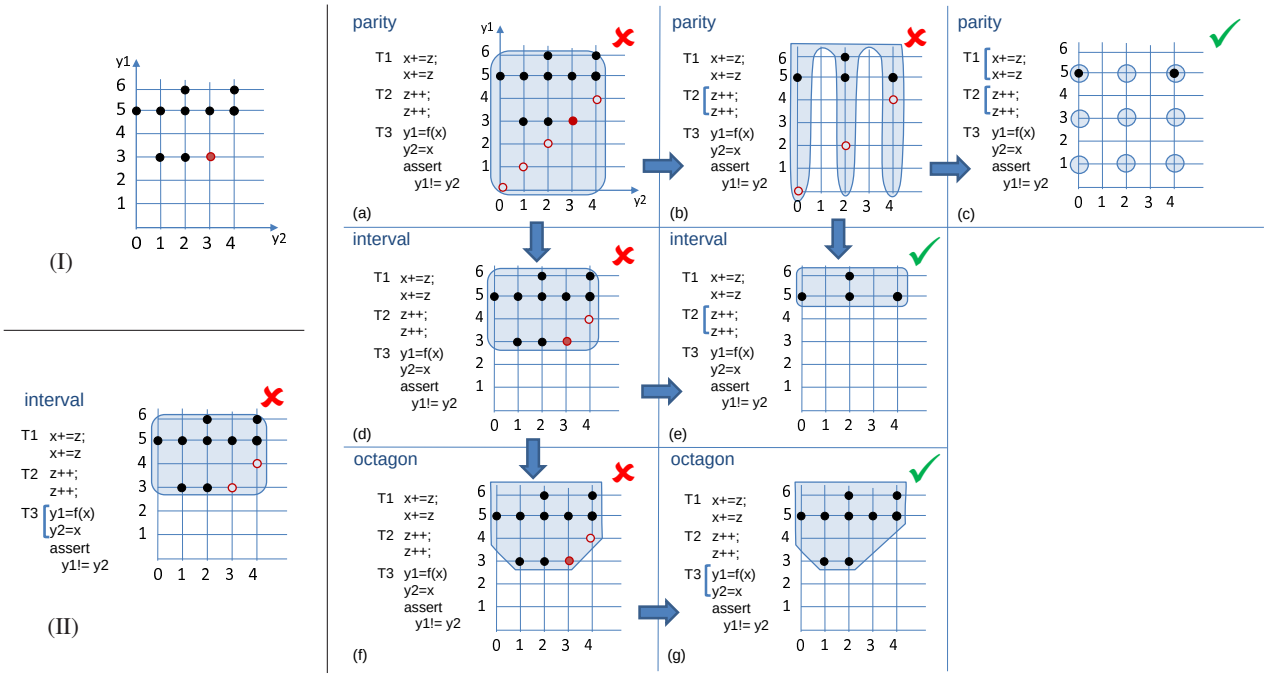

Figure 2. **(I)** Values of $y1$ and $y2$ that arise in the program of Fig. 1; **(II)** Atomic section around the assignments to $y1$ and $y2$ under interval abstraction; **(a-g)** Possible steps of the AGS algorithm: on each step, the algorithm can choose between refining the abstraction (down arrows) and modifying the program by avoiding certain interleavings (right arrows).

2.3 Inferring Synchronization under Parity Abstraction

We first show how the algorithm works using the parity abstraction over $y1$ and $y2$. The parity abstraction represents the actual value of a variable by its parity, and only observes whether the value is even or odd. Variables $y1$ and $y2$ take abstract values from $\{\bot, E, O, \top\}$, with the standard meaning.

The starting point, parity abstraction of the original program, is shown in Fig. 2 (a). It shows the concrete values of $y1$ and $y2$ that can arise during program execution, and their abstraction. The concrete values are shown as full circles and are the same as in Fig. 2 (I). Black circles denote the concrete values that satisfy the assertion, and red circle values that violate the assertion. The shaded area denotes the concretization of the abstract values computed for $y1$ and $y2$. The abstract values for both $y1$ and $y2$ are $\top$. As a result, the concretization (the shaded area) covers the entire plane. In particular, it covers concrete values that violate the assertion. Values that cannot arise in any concrete execution of the program (false alarms) are shown as hollow red circles in the figure.

The AGS algorithm performs abstract interpretation of the program from Fig. 1 using parity abstraction. In Fig. 3 we show part of the abstract transition system constructed by AGS. Fig. 3 only shows abstract states that can reach an error state. Error states are shown as dashed red line circles in the figure. The values of variables in a state are shown as a tuple $\langle pc_1, pc_2, pc_3, x, z, y1, y2\rangle$, where variables $y1$ and $y2$ take an abstract value from the parity domain. This transition system is very simple and in particular contains no cycles; however, this is only for illustrative purposes and the AGS algorithm handles all forms of abstract transition systems.

Under parity abstraction, there are several invalid interleavings. The choice which of them to eliminate first is important, as discussed in Section 5. The AGS algorithm first chooses to eliminate the invalid interleaving: $\pi_1 =$ `z++; x+=z; x+=z; z++; y1=f(x); y2=x; assert`. This interleaving is shown in Fig. 3 by emphasizing its edges (the right emphasized path in the figure). Under this interleaving, and under the parity abstraction, $y1 = \top$ and $y2 = \top$ (due to joins in the abstract transition system).

The AGS algorithm can now choose whether to try and eliminate this by either adding atomicity, or by refining the abstraction. Fig. 2 shows these alternatives, which we explain in detail in the rest of this section.

Eliminate π_1 by atomicity constraint: To eliminate this interleaving, the following constraint is generated: `[z++,z++]`. This step is shown as the step from Fig. 2 (a) to Fig. 2 (b). Note that the program in Fig. 2 (b) has an atomic section around the statements `z++` and `z++` in `T2`. This limits the concrete values that $y1$ and $y2$ can take, as shown by the full circles in Fig. 2 (b), compared to those on Fig. 2 (a). In particular, it eliminates the error state in which $y1$ and $y2$ both have the value 3 (no red full circle in the figure).

However, parity abstraction is not yet precise enough to verify the correctness of the resulting program, as shown by the shaded area in Fig. 2 (b). During abstract interpretation of the program, $y1$ takes both the values E and O, and thus goes to $\top$. The concretization (the shared area) therefore spans all possible concrete values of $y1$. The abstract value of $y2$ remains E, therefore the concretization (the shaded area) only contains even values of $y2$. The abstract values represent three points that violate the assertion, shown as hollow red circles in Fig. 2 (b).

After eliminating π_1 by adding the constraint `[z++,z++]`, the following (abstract) interleaving may violate the assertion: $\pi_2 =$ `x+=z;z++;z++;x+=z;y1=f(x); y2=x; assert`. This interleaving yields the abstract values $y1 = \top$ and $y2 = \top$ (due to joins), which may violate the assertion. The interleaving π_2 is shown in Fig. 3 as the left emphasized path in the figure.

Eliminate π_2 by atomicity constraint: To eliminate this interleaving, the following constraint is generated: `[x+=z,x+=z]`. This step is shown as the step from Fig. 2 (b) to Fig. 2 (c). The resulting overall constraint is: `[x+=z,x+=z]` $\wedge$ `[z++,z++]`

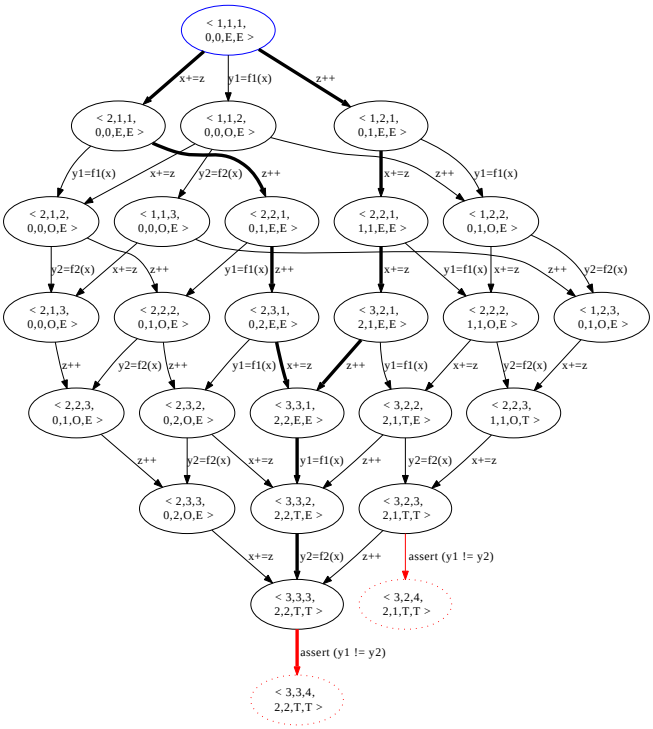

Figure 3. Partial abstract transition system for the program of Fig. 1. Only abstract states that can reach an error state are shown.

With this atomicity constraint, under the parity abstraction, there are no further invalid interleavings. This constraint is satisfied by a program that has the statements `x+=z` and `x+=z` of `T1` execute atomically, and the statements `z++` and `z++` of `T2` execute atomically. In this program, the abstract values are $y1 = O$ and $y2 = E$. These abstract values guarantee that the assertion is not violated, as shown in Fig. 2 (c).

Eliminate π_2 by abstraction refinement: After eliminating the interleaving π_1, all remaining concrete interleavings satisfy the assertion, but we could not prove it under parity abstraction. Instead of eliminating interleaving π_2 by adding atomicity constraints, as described above, we can choose to refine the abstraction from parity to interval, moving from Fig. 2 (b) to Fig. 2 (e). Interval abstraction is precise enough to prove this program.

2.4 Inferring Synchronization under Interval Abstraction

Instead of eliminating interleaving π_1 by adding an atomicity constraint, the algorithm can choose to try and eliminate π_1 by refining the abstraction from parity to interval. This corresponds to the step from Fig. 2 (a) to Fig. 2 (d). Under interval abstraction, the abstract values are $y1 = [3, 6]$ and $y2 = [0, 4]$, representing two points that may violate the assertion, as shown in figure Fig. 2 (d).

The algorithm can again choose to eliminate invalid interleavings by adding an atomicity constraint (step from Fig. 2 (d) to Fig. 2 (e)) or by abstraction refinement (step from Fig. 2 (d) to Fig. 2 (f)). In the former case, AGS produces the overall constraint:

$$\big(\texttt{[x+=z,x+=z]} \vee \texttt{[z++,z++]}\big) \\ \wedge \big(\texttt{[y1=f(x),y2=x]} \vee \texttt{[x+=z,x+=z]} \vee \texttt{[z++,z++]}\big)$$

This constraint requires only one of `T1` and `T2` to execute atomically. Fig. 2 (e) shows a program corresponding to one of the solutions, in which `T2` is atomic.

As apparent from the constraint above, `[y1=f(x),y2=x]` is not sufficient for showing the correctness of the program under the interval abstraction. The result of applying interval abstraction to the program implemented from this constraint is shown in Fig. 2 (II).

2.5 Inferring Synchronization under Octagon Abstraction

Finally, the octagon abstract domain [18], maintains enough information to only require atomicity as in the case with full information. In particular, it is sufficient to make `y1=f(x)` and `y2=x` execute atomically for the program to be successfully verified under Octagon abstraction, as shown in Fig. 2 (g).

3. Preliminaries

Transition System A transition system ts is a tuple $\langle \Sigma, T, \mathit{Init} \rangle$ where Σ is a set of states, $T \subseteq \Sigma \times \Sigma$ is a set of transitions between states, and $\mathit{Init} \subseteq \Sigma$ are the initial states. For a transition $t \in T$, we use $src(t)$ to denote the source state of t, and $dst(t)$ to denote its destination state.

For a transition system ts, a trace π is a (possibly infinite) sequence of transitions $\pi_0, \pi_1, \ldots$ such that for every $i > 0$, $\pi_i \in T$ and $dst(\pi_{i-1}) = src(\pi_i)$. For a finite trace π, $|\pi|$ denotes its length (number of transitions). We use $t.\pi$ to denote the trace created by concatenation of a transition t and a trace π, when $dst(t) = src(\pi_0)$.

A complete trace π is a trace that starts from an initial state: $src(\pi_0) \in \mathit{Init}$. We use $[\![ts]\!]$ to denote the (prefix-closed) set of complete traces of transition system ts.

Program Syntax We consider programs written in a simple programming language with assignment, non-deterministic choice, conditional goto, sequential composition, parallel composition, and atomic sections. The language forbids dynamic allocation of threads, nested atomic sections, and parallel composition inside an atomic section. Note that a program can be statically associated with the maximal number of threads it may create in any execution. Assignments and conditional goto statements are executed atomically. All statements have unique labels. For a program label l, we use $\mathit{stmt}(l)$ to denote the unique statement at label l.

We use *Var* to denote the set of (shared) program variables. To simplify the exposition, we do not include local variables in definitions, although we do use local variables in examples. There is nothing in our approach that prevents us from using local variables, but having local variables makes the formal definitions cumbersome. We assume that all program variables have integer values, initialized to 0.

Program Semantics Let P be a program with variables *Var*. Let k be the maximal number of threads in P, with thread identifiers $1, \ldots, k$. A state s is a triplet $\langle val_s, pc_s \rangle$ where $val_s \colon \mathit{Var} \to \mathit{Int}$ is a valuation of the variables, and $pc_s \colon \{1, \ldots, k\} \to \mathit{Int}$ is the program counter of each thread, which ranges over program labels in the code executed by the thread.

We define a transition system for a program P to be $\langle \Sigma_P, T_P, \mathit{Init}_P \rangle$, where transitions T_P are labeled by program statements. For a transition $t \in T_P$, we use $\mathit{stmt}(t)$ to denote the corresponding statement. We use $lbl(t)$ and $tid(t)$ to denote (unique) program label and thread identifier that correspond to $\mathit{stmt}(t)$, respectively.

A transition t is in T_P if all of the following conditions hold:

(a) the program counter of the thread $tid(t)$ in state $src(t)$ is at program label $lbl(t)$,
(b) the execution of the statement $\mathit{stmt}(t)$ from state $src(t)$ by thread $tid(t)$ results in state $dst(t)$,
(c) no other thread is inside an atomic section in state $src(t)$.

We use $[\![P]\!]$ to denote the set of traces of P, i.e., $[\![P]\!] = [\![ts]\!]$ where $ts = \langle \Sigma_P, T_P, \mathit{Init}_P \rangle$.

Abstraction Our method is based on abstract interpretation [10]. In this section, we quickly review relevant terminology that will be used throughout the paper.

An abstract domain is a complete join semilattice $A = \langle A, \sqsubseteq, \sqcup, \bot \rangle$, i.e., a set A equipped with partial order $\sqsubseteq$, such that for every subset X of A, A contains a least upper bound (or join), denoted $\sqcup X$. The bottom element $\bot \in A$ is $\sqcup \emptyset$. We use $x \sqcup y$ as a shorthand for $\sqcup\{x, y\}$.

In this paper, we assume that the abstract domain A is a powerset of abstract states, with (partially) disjunctive join. An abstract state s is ranging over an abstract domain $B = \langle B, \sqsubseteq_B, \sqcup_B, \bot_B \rangle$.

For $X \subseteq \Sigma_P$, the abstraction function α is defined by $\alpha(X) \stackrel{\text{def}}{=} \sqcup\{\beta(s) \mid s \in X\}$, where β is the abstraction function for the underlying domain of abstract states. For a given β, the abstraction α can vary anywhere on the range between "relational" and "cartesian", depending on the definition of join.

An abstract transformer for a program statement st is denoted by $[\![st]\!]_\alpha \colon A \to A$. For $a \in A$, the abstract transformer is defined pointwise: $[\![st]\!]_\alpha(a) \stackrel{\text{def}}{=} \sqcup\{[\![st]\!]_\beta(\sigma) \mid \sigma \in a\}$, where $[\![st]\!]_\beta$ is the abstract transformer for the underlying domain of abstract states.

We abuse the notation slightly and use α to collectively name all the components of an abstract interpreter: its abstract domain, including the underlying domain of abstract states, abstract transformers, and widening operator, if defined.

We define an abstract transition system for P and α to be $\langle \Sigma_P^\natural, T_P^\natural, \mathit{Init}_P^\natural \rangle$, where $\mathit{Init}_P^\natural = \alpha(\mathit{Init}_P)$, and a transition (σ, σ') labeled by a program statement st is in $T_P^\natural$ if and only if $[\![st]\!]_\beta(\sigma) \sqsubseteq_B \sigma'$.

We use $[\![P]\!]_\alpha$ to denote the set of abstract traces of P, i.e., $[\![P]\!]_\alpha = [\![ts]\!]$ where ts is the abstract transition system for P and α, in which $\Sigma_P^\natural$ is the result of abstract interpretation, i.e., the set of abstract states at fixed point.

Specification The user can specify a state property S, which describes a set of program states. This property can refer to program variables and to the program counter of each thread (e.g., to model local assertions). Our approach can be extended to handle any temporal safety specifications, expressed as a property automaton, by computing the synchronous product of program's transition system and the property automaton [9].

Given a (concrete or abstract) state s, we use $s \models S$ to denote that the state s satisfies the specification S. We lift it to traces as follows. A trace π satisfies S, denoted by $\pi \models S$, if and only if $src(\pi_0) \models S$ and for all $i \geq 0$, $dst(\pi_i) \models S$. A set Π of (concrete or abstract) traces satisfies S, denoted by $\Pi \models S$ if and only if all traces in it satisfy S.

4. Computing a Safe Schedule Under Abstraction

Algorithm 1 provides a declarative description of abstraction-guided synthesis. The algorithm takes an input program, a specification, and an abstraction, and produces a (possibly modified) program that satisfies the specification.

The main loop of the algorithm selects an abstract trace π of the program P such that π satisfies the atomicity formula φ, but does not satisfy the specification S. Then, the algorithm attempts to eliminate this invalid interleaving π by either:

- adding atomicity constraints: the procedure `avoid` generates atomicity constraints that disable π. The constraints generated by `avoid` for π are accumulated by AGS in the formula φ.
- refining the abstraction: using a standard abstraction refinement approach (e.g., [8, 3]) to refine the abstraction.

On every iteration, the loop condition takes into account the updated φ and α when choosing an invalid interleaving π.

Some of the (abstract) invalid interleavings may correspond to concrete invalid interleavings, while others may be artifacts of the abstraction. The choice of whether to eliminate an interleaving via abstraction refinement, or by adding atomic sections, is left as non-deterministic choice (denoted by $\star$ in the algorithm). In this section, we assume that it makes the right choices (for example, only picks refinement when it is indeed possible to eliminate π using refinement). In Section 5, we discuss how to implement it.

When all invalid interleavings have been eliminated, AGS calls the procedure `implement` to find a solution for the constraints accumulated in φ.

Algorithm 1: Abstraction-Guided Synthesis.

Input: Program P, Specification S, Abstraction α
Output: Program satisfying S under α

```
1  φ = true
2  while true do
3      Π = {π | π ∈ [[P]]_α ∩ [[φ]], π ⊭ S}
4      if Π is empty then return implement(P,φ)
5      π = select trace from Π
6      if shouldAvoid(π, α) then
7          ψ = avoid(π)
8          if ψ ≠ false then φ = φ ∧ ψ
9          else abort
10     else
11         α' = refine(α, π)
12         if α' ≠ α then α = α'
13         else abort
14     end
15 end
```

Function `avoid(π)`

Input: Trace π
Output: Atomicity constraint for avoiding π

```
ρ = false
foreach i = 0, ..., |π| do
    if exists j > i + 1 such that tid(π_i) = tid(π_j) and
    for all l such that i < l < j, tid(π_i) ≠ tid(π_l)
    then ρ = ρ ∨ [lbl(π_i), lbl(π_j)]
end
return ρ
```

Function `implement(P, φ)`

Input: Program P, atomicity formula φ
Output: Program with atomic sections satisfying φ

```
Find a minimal satisfying assignment Γ ⊨ φ
P' = P with adding atomic sections in atomize(Γ)
return P'
```

4.1 Generating Atomicity Constraints

The procedure `avoid` takes a trace π as input, and generates an atomicity constraint that describes all context switches in π, and thus describes all possible ways to eliminate π by adding atomic sections to the original program.

The atomicity constraint generated by `avoid` is a disjunction of atomicity predicates. An atomicity predicate requires that a pair of consecutive program statements execute atomically, without interleaving execution of other threads between them.

Formally, given a program P, and a pair of program labels l and l', we use $[l, l']$ to denote an *atomicity predicate*. In our examples, we write $[\mathit{stmt}(l), \mathit{stmt}(l')]$ instead of $[l, l']$. An *atomicity formula* is a conjunction of disjunctions of atomicity predicates.

Let π be a trace in a (concrete or abstract) transition system of P. We say that π satisfies $[l, l']$, denoted by $\pi \models [l, l']$, if and only if for all $0 \leq i$, if $lbl(t_i) = l$ and $i + 1 < |\pi|$, then $lbl(t_{i+1}) = l'$ and $tid(t_i) = tid(t_{i+1})$.

A set of traces Π satisfies an atomicity predicate p, denoted by $\Pi \models p$, if and only if all the traces in Π satisfy p. Similarly, we interpret conjunctions and disjunctions of atomicity predicates as intersection and union of sets of traces. The set of traces that satisfy an atomicity formula φ is denoted by $[\![\varphi]\!]$.

The procedure `avoid` only generates atomicity predicates for neighboring locations (locations that appear in the same thread, where one location immediately follows the other), with the intuitive meaning that no operation is allowed to interleave between the execution of these neighboring locations.

The algorithm identifies all context switches in π as follows. A context switch after transition π_i occurs if there is another transition π_j by the same thread later in the trace, but not immediately after π_i. Then, if the transition π_j is the first such transition after π_i, we generate the atomicity predicate $[lbl(\pi_i), lbl(\pi_j)]$.

In the case of an invalid sequential interleaving, an interleaving in which each thread runs to completion before it context-switches to another thread, it is (obviously) impossible to avoid the interleaving by adding atomic sections. In such cases, `avoid` returns $false$ and AGS aborts.

4.2 Implementing Atomicity Constraints

The procedure `implement` takes a program P and an atomicity formula φ as input. An atomicity formula can be seen as a formula in propositional-logic, where the atomicity predicates are treated as propositional (boolean) variables. Note that the atomicity formula is in positive CNF, and thus it is always satisfiable.

The procedure constructs a program P' by finding a minimal satisfying assignment for φ, i.e., a satisfying assignment with the smallest number of propositional variables set to *true*. The atomicity predicates assigned to true in that assignment are then implemented as atomic sections in the program.

Our approach separates the characterization of valid solutions from their implementation. The atomicity formula φ maintained in the algorithm provides a symbolic description of possible solutions. In this paper, we choose to realize these by changing the program and adding atomic sections. However, these could be realized using other synchronization mechanisms, as well as by controlling the scheduler of the runtime environment (if such scheduler exists).

In general, there could be multiple satisfying assignments for φ, corresponding to different additions of atomic sections to the input program P. Usually, we are interested in minimal satisfying assignments, as they represent solutions that do not impose redundant atomic sections.

To realize a satisfying assignment $\Gamma \models \varphi$ as atomic sections, we define $atomize(\Gamma)$ to extract the minimal (contiguous) atomic sections from the assignment. Towards this end, we construct the set of program labels in which context switches are not permitted by Γ: $L = \{l' \mid [l, l'] \in \Gamma\}$. For every maximally-connected component of L in the control-flow-graph of the original program, we find the immediate dominator and postdominator, and add (begin and end) atomic section at these labels, respectively. This may cause extra statements included in an atomic section, eliminating additional interleavings. This situation is sometimes unavoidable when implementing atomicity constraints using atomic sections.

It is possible that implementing an assignment Γ results in eliminating additional interleavings even when there are no extra statements in the atomic section. Consider the example of Fig. 4. In this example, `T2` cannot interleave with the first iteration of the loop in `T1`. But once the first iteration is over, it can interleave with any other iteration. However, since we require implementation via atomic sections, the only implementable solution is to add an atomic section around the statements `x++` and `x++` inside the loop, forcing every iteration of the loop to be executed atomically.

4.3 Abstraction Refinement

The procedure `refine` takes an interleaving π as input and attempts to refine the abstraction in order to avoid π. For that to be possible, π has to be an artifact of the abstraction, and not correspond to a concrete invalid interleaving. AGS tries to refine the abstraction by calling `refine`, but if the abstraction cannot be refined, and `refine` returns the same abstraction, AGS aborts.

In this paper, we focus on the procedure for restricting invalid interleavings, and can leverage any standard refinement scheme (e.g., [8, 3, 4, 22]). In the examples, we use two kinds of simple refinements: one that moves to another abstract domain (Section 2), and one that varies the set of variables that are abstracted relationally (Section 7).

4.4 Choosing Interleaving π to Eliminate

Since our program modifications consist of adding atomic sections, we cannot eliminate sequential executions (which have no context switches). It is therefore required that we can verify the correctness of the sequential runs of the program under the given abstraction.

In fact, for verifying the correctness of interleavings that involve fewer context-switches, less precise abstractions can be sufficient.

Generally, it is natural to consider interleavings in an increasing order of the number of context switches. Atomicity constraints obtained for interleavings with a lower number of context switches restrict the space that needs to be explored for interleavings with higher number of context switches.

4.5 Program Modification vs. Abstraction Refinement

When an invalid interleaving π is detected, a choice has to be made between refining the abstraction and adding an atomicity constraint that eliminates π. This choice is denoted by the condition $shouldAvoid(\pi, \alpha)$ in the algorithm. Apart from clear boundary conditions outlined below, this choice depends on the particular abstractions with which the algorithm is used.

When π is a sequential interleaving, and `avoid` is realized as the addition of atomic sections, it is impossible to add atomicity constraints to avoid π. Therefore, in this case, the only choice is to refine the abstraction (if possible). Hence, the condition $shouldAvoid(\pi, \alpha)$ is set to return $false$ when π is a sequential interleaving.

Similarly, depending on the refinement framework used, it may be impossible to further refine the abstraction α. For example, when using a fixed sequence of abstraction with increasing precision (as in Section 2), upon reaching the most precise abstraction in the sequence, there's no way to further refine the abstraction. Therefore, in this case, the only choice is trying to avoid the interleaving π, and the condition $shouldAvoid(\pi, \alpha)$ returns $true$ when it is known a priori that α cannot be refined anymore.

For refinement schemes that use symbolic backwards execution to find a concrete counterexample (e.g., [8, 3]), the condition $shouldAvoid(\pi, \alpha)$ can be based on the result of the symbolic execution. When the refinement scheme is able to find a concrete counterexample, $shouldAvoid(\pi, \alpha)$ can choose to repair, using the concrete counterexample as basis. If the refinement scheme fails to find a concrete counterexample, but also fails to find a spurious path for abstraction refinement, $shouldAvoid(\pi, \alpha)$ can again choose to repair, as refinement cannot be applied.

Attempting verification with a refined abstraction may fail due to state explosion. In most cases there is no way to check for such failure a priori in the condition $shouldAvoid(\pi, \alpha)$. Practically, it is useful to invoke the verification procedure as a separate task, and implement a backtracking mechanism for the refinement when verification fails to terminate after a certain time. Backtracking the refinement may enable successful verification of a more constrained variant of the program.

```
T1 {                    |  T2 {
 while (*) {            |   if (x==1) {
  x++                   |    assert false
  x++                   |   }
 }                      |  }
}
```

Figure 4. Limitations of implementability. Correctness only requires the first iteration of the loop in `T1` be executed atomically. Implementability forces every iteration to be executed atomically.

5. Abstraction Guided Synthesis

In the previous section, we described the AGS algorithm in a declarative manner, and omitted some details that we now address:

- how do we compute $[\![P]\!]_\alpha$?
- how do we obtain an interleaving $\pi \in [\![P]\!]_\alpha \cap [\![\varphi]\!]$ and $\pi \not\models S$?
- how do we choose, on every step of the algorithm, whether to add atomicity constraints or to refine the abstraction?

To realize Algorithm 1, we first use standard abstract interpretation to compute the set of abstract states $\Sigma_P^\natural$ reachable from $Init_P^\natural$ under abstraction α. Then, we explore the invalid interleavings and eliminate them. The algorithm is amenable to several optimizations, and we describe them later in this section.

In the pseudocode of Algorithm 1, we replace the declarative expression in Line 3 with a call to function `Traces`:

$$\Pi = \texttt{Traces}(Init_P^\natural, Bad_P^\natural, \Sigma_P^\natural, \varphi)$$

in Algorithm 1, where $Bad_P^\natural$ is the set of reachable error states: $\{\sigma \in \Sigma_P^\natural \mid \sigma \not\models S\}$.

Function `Traces` (X, Y, V, φ)

```
Input: Set of abstract states X, Y, V, Atomicity Formula φ
Output: Set of traces from X to Y passing in V, satisfying φ
workset = {t | src(t) ∈ V \ X, dst(t) ∈ Y}
result = {t | src(t) ∈ X, dst(t) ∈ Y}
while workset is not empty do
    π = select and remove interleaving from workset
    foreach Statement st and state σ ∈ V such that
    [[st]]_β(σ) ⊑_B src(π_0) do
        t = transition (σ, src(π_0)) labeled with st
        π' = t.π
        if π' ⊨ φ and π' is acyclic then
            if σ ∈ X then result = result ∪ {π'}
            else workset = workset ∪ {π'}
        end
    end
end
return result
```

The function $\texttt{Traces}(X, Y, V, \varphi)$ enumerates all traces that start in a state in $X \subseteq V$, end in a state in $Y \subseteq V$, go only through states in V and satisfy the atomicity constraint φ. It works by performing a backward exploration starting from states in Y and extending interleaving suffixes backwards. A suffix is further extended only as long as it satisfies φ. Thus, the algorithm leverages the constraints that are already accumulated in the atomicity formula φ to prune the interleavings that have to be explored. The use of φ is critical for the practicality of the approach, as shown experimentally in Section 7.

Exploring φ-enabled statements We say that statement st is *φ-enabled in state σ* when executing st from σ does not contradict φ. Formally, given a set of states V, the condition $\texttt{enabled}(st, \sigma, \varphi, V)$ holds if and only if for every pair of transitions t and t' such that $src(t) \in V$, $dst(t') \in V$, $dst(t) = src(t') = \sigma \in V$ and $stmt(t') = st$, the partial trace $t.t'$ satisfies φ.

For example, if φ is `[a,c]` then st is not φ-enabled in σ, in the partial state space shown on the right. However, if φ is `[a,c]` $\vee$ `[b,c]`, then st is φ-enabled in σ.

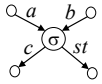

Algorithm 2: Abstraction-Guided Synthesis.

```
   Input: Program P, Specification S, Abstraction α
   Output: Program satisfying S under α
1  states = workset = Init_P^♮
2  φ = true
3  while workset is not empty do
4      σ = select and remove state from workset
5      foreach Statement st do
6          if enabled(st, σ, φ, states) then
7              σ' = [[st]]_β(σ)
8              if σ' ⊭ S then
9                  select
                   π ∈ Traces(Init_P^♮, {σ'}, states\workset, φ)
10                 if shouldAvoid(π, α) then
11                     ψ = avoid(π)
12                     if ψ ≠ false then
13                         φ = φ ∧ ψ
14                         states = workset = Init_P^♮
15                         disabled = ∅
16                     else abort
17                 else
18                     // refine(π)
19                 end
20             else
21                 if {σ'} ⋢ states then
22                     states = states ⊔ {σ'}
23                     X = {σ'' ∈ states | σ' ⊑_B σ''}
24                     workset = workset ⊔ X
25                 end
26             end
27         end
28     end
29 end
30 return implement(P,φ)
```

Forward Pruning using φ Algorithm 2 is an optimized version of Algorithm 1. In the optimized algorithm, we focus on the exploration code, and on the code for avoiding an interleaving (Lines 11-16), the code for refinement is similar and is abbreviated to a comment in Line 18. The algorithm combines (forward) abstract interpretation of the program, with (backward) exploration of invalid interleavings. The main idea of the algorithm is to use the constraints accumulated in φ to restrict the space that has to be explored both forward and backward. In particular, this optimization avoids constructing the entire (unrestricted) transition system upfront.

The abstract interpretation part of the algorithm is standard, and uses a workset to maintain abstract states that should be explored. Once the workset is empty we know a fixed point is reached.

At every point, forward exploration of new states is restricted by the current constraints accumulated in φ (Line 6). For every invalid interleaving π, the formula φ represents all the possible ways to eliminate π. This means that the algorithm only restricts further exploration when the next exploration step contradicts *all possible ways* to eliminate existing invalid interleavings.

In the algorithm, we use the join operator of the abstract domain to add new states to the set $states$ of explored abstract states (Line 24). More generally, the algorithm can use a widening operator [10] when required. To determine whether a state should be added to the set of states, we check whether the state is already represented in $states$ (Line 21).

Rebuilding Parts of the Transition System Instead of rebuilding the whole transition system whenever we add a constraint to φ (Line 14), or whenever we refine the abstraction, we can rebuild only the parts of the transition system that depend on the modification. Following approaches such as [13], we can invalidate only the parts of the abstract transition system that may be affected by the refinement, and avoid recomputation of other parts.

Lazy Abstraction Algorithm 2 need not maintain the same abstraction across different interleavings. The algorithm can be adapted to use lazy abstraction refinement as in [13]. Instead of maintaining a single homogenous abstraction α for the entire program, we can maintain different abstractions for different parts of the program, and perform lazy refinement.

Simplification of φ Rather than taking the conjunction of constraints as they are accumulated in φ, we preform (propositional) simplification of φ on-the-fly. This is required in practice, as the number of terms added to φ may be large even for small programs.

Multiple Solutions The algorithm as described here only yields a single minimal solution. In practice (and in our implementation, described in Section 7), it is often desirable to present the user with a range of possible solutions and let the user make her own choice.

6. Correctness and Minimality

In this section, we show that Algorithm 1 computes a correct program with smallest atomic sections, assuming the abstraction is fixed. At the end, we discuss the effect of abstraction refinement, and correctness of Algorithm 2.

6.1 Correctness

In this section, we assume that the abstraction is fixed, i.e., $shouldAvoid$ in Algorithm 1 always returns $true$. The following theorem says that a run of the AGS algorithm terminates with either an abort or a valid program.

THEOREM 6.1 (Correctness). *A run of the AGS algorithm terminates with either an abort or returns a program P' such that*
(1) P' satisfies S under α, i.e., $[\![P']\!]_\alpha \models S$, and
(2) P' admits a subset of interleavings of the original program P, i.e., $[\![P']\!]_\alpha \subseteq [\![P]\!]_\alpha$.

Sketch of Proof: In every iteration, the AGS algorithm eliminates at least one simple path to error state from the abstract transition system. As a result, the abstract transition system may be modified to take into account the updated atomicity formula φ. However, the abstract transition system is always modified in a way that does not introduce any new paths, in particular it has no new paths to error states.

Because the number of simple paths is finite, the `while` loop in the AGS algorithm terminates either by eliminating all simple paths to error, or finding a path to error that has no context switches and thus it cannot be eliminated by our method. In the latter case, the algorithm aborts. In the former case, the set of traces Π is empty. That is, any execution of P that respects the atomicity formula φ satisfies S under the abstraction α. Let P' be the program returned by `implement(P,`φ`)` in this case. The interleavings of P' is a subset of those of P permitted by φ under α: $[\![P']\!]_\alpha \subseteq [\![P]\!]_\alpha \cap [\![\varphi]\!]$. Therefore, P' satisfies S under α and AGS algorithm returns P'.

The AGS algorithm cannot fix a program whose sequential executions do not satisfy S under α. Otherwise, if there is a way to add atomic sections to P such that the result satisfies S under α, then there exists a run of the AGS algorithm that does not abort, and computes a result. In the worst case, it makes the program always execute sequentially.

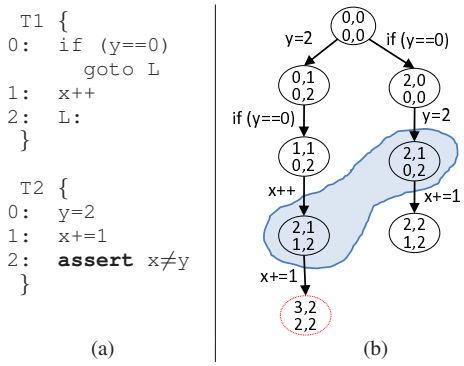

Figure 5. Example demonstrating the effect of join and the choice of different abstract traces to eliminate.

THEOREM 6.2. *If the sequential executions of P satisfy S under α, then there exists a run of AGS algorithm that does not abort.*

In Algorithm 2, at Line 9, we always choose from `Traces` a trace π that has context switches, if there is one. It guarantees that no run of AGS algorithm aborts if the sequential version of P is valid.

The toy example in Fig. 5(a) has a single invalid interleaving: `y=2;if(y==0);x++;x+=1;assert`, as shown on Fig. 5(b). However, under parity abstraction, there are two invalid interleavings, due to join (shaded area). One of them is the abstraction of the concrete invalid interleaving, denoted by π_1. The other one is a sequential interleaving, denoted by π_2, in which `T1` executes first, and then `T2`. If the AGS algorithm first chooses to eliminate π_2, it will abort, because there are no context switches to disable. However, if we chose π_1 first, `avoid` will return the atomicity constraint `[y=2,x+=1]`, and the program will be successfully verified under this constraint, using parity abstraction. Similarly, we can construct an example in which a wrong choice leads to larger atomic sections than necessary.

6.2 Minimality

Next, we define the notion of a minimally-atomic program, and show how to use the AGS algorithm to compute all minimally-atomic programs for a given input program, specification and abstraction.

Let Γ be a set of atomic predicates that refer to a program P. In AGS algorithm, we obtain Γ as a satisfying assignment to the atomicity formula φ. Recall from Section 4.2 that there is a unique way to realize Γ by adding atomic sections to P. Let us denote the resulting program by $P\,|_\Gamma$.

Let P' be obtained from P by adding atomic sections. We use $\Gamma(P, P')$ to denote the (unique) set of atomic predicates that corresponds to these atomic sections.

Minimally Atomic Programs A valid program is minimally-atomic when removing or shrinking any atomic section in it makes it invalid.

DEFINITION 6.3 (Minimally Atomic). *Consider a program P and an abstraction α. Let P' be a program obtained from P by adding atomic sections. P' is minimally-atomic with respect to α if and only if $[\![P']\!]_\alpha \models S$ and for every program P'' obtained from P by adding atomic sections, if $\Gamma(P, P'') \subset \Gamma(P, P')$, then $[\![P'']\!]_\alpha \not\models S$.*

The condition $\Gamma(P, P'') \subset \Gamma(P, P')$ means that the atomic sections of P'' is a (strict) subset of those of P'.

We use $MA(P, \alpha)$ to denote the set of all minimally-atomic programs with respect to α that can be obtained from P. The

programs in $MA(P, \alpha)$ have incomparable sets of atomic sections, i.e., for every pair $P', P'' \in MA(P, \alpha)$, $\Gamma(P, P') \not\subset \Gamma(P, P'')$. However, they may have the same set of traces under α (and even concrete traces). When the abstraction α is not precise enough to prove that all sequential executions of P satisfy S, $MA(P, \alpha)$ is empty. In the rest of this section, we show that every minimally-atomic program can be implemented by AGS algorithm.

THEOREM 6.4 (Minimality). *For every minimally-atomic program $P' \in MA(P, \alpha)$, there exists a run of the AGS algorithm that returns P'.*

Sketch of Proof: Let $P' \in MA(P, \alpha)$ and let $\Gamma = \Gamma(P, P')$. We show that Γ is a satisfying assignment to φ computed in some run of AGS algorithm, i.e., for some sequence of invalid interleavings picked by AGS to be eliminated.

Let Γ' be a maximal subset of Γ such that $P \mid_{\Gamma'} \neq P'$. Because P' is a minimally-atomic program w.r.t. α, $[\![P \mid_{\Gamma'}]\!]_\alpha \not\models S$. There exists an atomicity predicate p such that $[\![P \mid_{\Gamma' \cup \{p\}}]\!]_\alpha \models S$. Thus, there is an invalid interleaving π in $[\![P \mid_{\Gamma'}]\!]_\alpha$ that is eliminated by p. Note that the atomicity predicates in $\Gamma \setminus (\Gamma' \cup \{p\})$ are the result of `atomize`. That is, $P' = P \mid_{\Gamma =} P \mid_{\Gamma' \cup \{p\}}$, because P' is the result of `implement(P,`φ`)` which chooses $\Gamma' \cup \{p\}$ as the minimal assignment it implements.

Assume that the invalid interleaving π is picked by AGS in the last iteration. Atomicity constraint ψ generated by `avoid(`π`)` will include p as one of its disjuncts. Suppose that there exist a run of AGS that produces φ' such that Γ' is a minimal satisfying assignment for φ'. Then, $\Gamma' \cup \{p\}$ is a minimal satisfying assignment to $\varphi = \varphi' \wedge \psi$ and φ is produced in the last iteration of AGS.

Similarly, we can continue subtracting atomicity predicates from Γ', constructing the sequence of invalid interleavings backwards, until we run out of atomicity predicates.

The following proposition is much stronger than Theorem 6.4, as it requires that a single run of the AGS algorithm yield all minimally-atomic programs. Moreover, the minimally-atomic programs exactly correspond to all minimal satisfying assignments of the atomicity formula φ computed by that run.

PROPOSITION 6.5 (Minimality-Strong). *If the sequential executions of P satisfy S under α, then there exists a run of the AGS algorithm that yields atomicity formula φ such that*

- *for every minimal satisfying assignment A to φ, the program `implement`$(P, A) \in MA(P, \alpha)$,*
- *for every $P' \in MA(P, \alpha)$, there exists a minimal satisfying assignment A for φ, such that `implement`(P, A) returns P'.*

Correctness and Minimality of Algorithm 2 Correctness of the operational version of the AGS algorithm, given in Algorithm 2, follows from the fact that an invalid interleaving eliminated by Algorithm 2 from a partial transition system is also an invalid interleaving that can be chosen by an iteration of Algorithm 1 from the corresponding full transition system. Minimality follows from the fact that the order of (forward) exploration in Algorithm 2 can be chosen to discover error states in a way that exhibits any sequence of minimal invalid interleavings.

Abstraction Refinement If refinement is not guaranteed to terminate, then AGS algorithm is not guaranteed to terminate. The reason is that every refinement step may produce new simple invalid interleavings. When the refinement is guaranteed to be monotonic, i.e., abstraction is more precise in every step (e.g., parity to intervals is not monotonic), we can attain minimality under abstraction refinement, by discarding the atomicity constraints φ after each refinement step. When the refinement is not monotonic, we can define a minimally-atomic program to respect any of the explored abstractions. In the case of lazy abstraction, which refines only part of the state-space, the definition of minimality is even more involved.

Finding a minimally-atomic program requires backtracking and it is at least exponential in the size of the abstract transition system of the input program, inline with the known complexity bounds for game-based synthesis [14]. Thus, it is more valuable to invest into a good heuristic. The simple heuristics that we use in the AGS algorithm produce reasonable, and often minimal, synchronization in practice, as we show in the next section.

Program	Refinement Steps	Avoid Steps
Double Buffering	1	2
Defragmentation	1	8
3D Update	2	23
Array Removal	1	17
Array Init	1	56

Table 1. Experimental Results.

7. Experience

We built a prototype tool named GUARDIAN based on the AGS algorithm of Section 5. We applied GUARDIAN to several interesting programs, inspired by real applications, which we describe next. The abstractions we used are variants of parity and interval domains, where the abstractions differ in what variables are kept relational.

Table 1 summarizes our experimental results. Note that all of our example programs are infinite state, and hence require abstraction for full verification. In our experiments, we were interested in exploring the space of fixes under several abstractions. Even when GUARDIAN found a solution with the original abstraction, we still let it explore solutions with finer abstractions. For every program in the table, we report the number of refinement steps, and number of `avoid` steps performed by the algorithm. In Table 2, we report the atomicity constraints found by GUARDIAN for programs whose code is shown in the paper (atomicity constraints refer to the code).

When using φ-pruning, all experiments ran in less than 10 minutes. Without using φ to restrict exploration, most programs went out of memory exploring a hopelessly large (and redundant) space of interleavings. To enumerate minimal assignments for the atomicity constraints constructed by our algorithm, GUARDIAN uses a model enumerator [1]. In the rest of this section, we describe some of our examples in more detail.

7.1 Abstract Domain

In our examples, the abstract domain is a powerset of abstract states. An abstract state s is a tuple $\langle val^\natural_s, pc_s \rangle$, where $val^\natural$ maps program variables to their abstract values, ranging over an abstract domain such as parity, sign or interval.

For $X \subseteq \Sigma_P$, the abstraction function α is defined as follows:

$$\alpha(X) \stackrel{\text{def}}{=} \sqcup\{\langle \theta(val_s), pc_s \rangle \mid s \in X\}$$

where θ maps concrete values of program variables to their abstract values. We use $\sqcup_\theta$ and $\sqsubseteq_\theta$ to denote the join and order of abstract values. The values of program counters are preserved.

We use the following join and partial order, parametric on the variables tracked relationally. Let $V \subseteq Var$ be a subset of variables. The join $\sqcup$ for the abstract domain A is defined using a relational join over a subset of variables in V and cartesian join over the rest of the variables:

$$s_1 \sqcup s_2 \stackrel{\text{def}}{=} \begin{cases} \{s_1, s_2\} & \text{if } pc_{s_1} \neq pc_{s_2} \text{ or} \\ & \text{exists } v \in V \text{ s.t. } val^\natural_{s_1}(v) \neq val^\natural_{s_2}(v) \\ \{\langle val^\natural_{s_1} \sqcup_\theta val^\natural_{s_2}, pc_{s_1} \rangle\} & \text{otherwise} \end{cases}$$

For $V = Var$ this defines relational join. For $V = \emptyset$ this defines cartesian join. Most of our examples vary the abstraction by varying

Program	Abstraction (set of tracked variables V)	Solution (atomicity constraints)
DBuffer	$\emptyset$	[fill:L1,fill:L2]) $\vee$ ([render:L1,render:L2]
	$\{Fill, Render\}$	*true*
Defrag	$\{Barrier, Region, F1, F2, empty\}$	[D:L1,D:L2] $\wedge$ [U:L1,U:L2] $\wedge$ [U:L2,U:L3] $\wedge$ [U:L3,U:L4]
	$\{Barrier, Region, F1, F2, empty, i, j\}$	[D:L1,D:L2] $\wedge$ [U:L1,U:L2]
3D update	$\{x2, x3, y3, z1, z3\}$	[P1:L2,P1:L3] $\wedge$ [P2:L2,P2:L3] $\wedge$ [P1:L8,P1:L9]
	$\{x2, x3, y3, z1, z3, y2, z2\}$	[P1:L2,P1:L3] $\wedge$ [P2:L2,P2:L3])
	$\{x2, x3, y3, z1, z3, y2, z2, x1, y1\}$	*true*

Table 2. Abstraction and solutions for some of the example programs

```
int Fill = 1;
int Render = 0;
int i = j = 0;
fill() {
 L1:if (i < N) {
 L2:  Im[Fill][i] = read();
 L3:  i += 1;
 L4:  goto L1;
   }
 L5: Fill ^= 1;
 L6: Render ^= 1;
 L7: i = 0;
 L8: goto L1;
}
```

```
render() {
 L1:if (j < N) {
 L2:  write(Im[Render][j]);
 L3:  j += 1;
 L4:  goto L1;
   }

 L5: j = 0;
 L6: goto 1;
}
main() {
 fill() || render();
}
```

Figure 6. Double Buffering

the relationality in the join. Because the join is parametric on the set V, in the presentation of our examples, we only vary the value of V. The value of V for each example is shown in Table 2. The partial order $\sqsubseteq$ on A is defined as follows: for all $Y, Y' \subseteq A$, $Y \sqsubseteq Y'$ if and only if for all $s_1 \in Y$ there exists $s_2 \in Y'$ such that $pc_{s_1} = pc_{s_2}$, and $val^{\natural}_{s_1} \sqsubseteq_\theta val^{\natural}_{s_2}$.

7.2 Double Buffering

This example is motivated by the mechanism of double buffering. Variants of this mechanism are used in a variety of settings, from computer graphics to device drivers. This scheme is illustrated in Fig. 6. There are two buffers of images (`Im`) of length N. The filler process fills the buffer indexed by the variable `Fill`, while at the same time the rendering process consumes the buffer indexed by variable `Render`. When the filling completes, the values of the two variables are swaped and the filling restarts. The rendering process simply renders the image indexed by variable `Render`. To avoid clutter, we assume that rendering is at least twice as fast as filling and hence before `Render` is changed, the value of its buffer has been written to the screen at least once (writing to the screen is idempotent and hence can be repeated).

Specification: We would like to prove that the filler and renderer processes never access the same location simultaneously. Formally:

$$pc(fill) = L2 \wedge pc(render) = L2 \Rightarrow \neg(Fill = Render \wedge i = j)$$

Result: Our first solution is obtained with a cartesian parity abstraction. This abstraction loses relationship between variables `Fill`, `Render`, `i` and `j` when states are joined. Formally, the set V of tracked variables is empty ($V = \emptyset$). Recall that the program counters are always kept relational.

With a more refined abstraction, GUARDIAN proves the correctness of the original program. The key reason why we succeeded in this case is that this abstraction maintains the relationship between the values `Fill` and `Render` on each iteration of the loop and can show that these two variables are never equal. In this example, refining the abstraction led to proving the program without any necessary fixes. Further refining the abstraction (e.g. inserting variable i or j in the set V) is not necessary.

7.3 Concurrent Defragmentation

This example is inspired by the problem of defragmentation. Defragmentation algorithms are used in various storage management scenarios (e.g., memory, disk storage) to increase space utilization. In many cases, defragmentation takes place concurrently with an executing application.

In Fig. 7, we show a simplified system where one process called `Defragment` performs memory compaction concurrently with another process called `Update` which allocates new entries in memory. The memory is organized as an array of entries called $Pages$. The size of the array N is unknown a-priori. Each entry in the array is either occupied (set to *true*) or free (set to *false*). In practice, an entry may correspond to a heap object or a file on a disk drive. Typically, each entry will also contain various other data fields, which we have omitted here for simplicity.

To avoid synchronization on each entry, the two processes should always work on disjoint regions of memory. To ensure that, at the start of their operation, the two processes handshake and then each picks a separate region to work with (labels L1-L2 in each process). Note that the handshake is not deterministic, and processes could select different regions on different handshakes. In our case, there are two regions, with the first region containing memory locations with an even index and the second region containing memory locations with an odd index.

`Defragment` works by iterating over the array and moving all used entries to one side of the page. `Update` works by selecting a memory location and updating it if some condition holds.

Specification: The processes should always access disjoint memory locations when at the program points accessing shared memory locations. (We omit the specification as it is long and tedious).

Result: The resulting constraints are shown in Table 2. The names of the processes have been abbreviated using their first letter. Note that the original program is incorrect (variable $Region$ is incremented without any synchronization), and with this more refined abstraction, the inferred correction is not a false positive (e.g. it is not due to an imprecision of the abstraction). However, the constraint [U:L2,U:L3] $\wedge$ [U:L3,U:L4] inferred with the coarser cartesian abstraction is due to the imprecision of the abstraction.

7.4 3D Grid Computation

Consider a concurrent program that updates values in a 3 dimensional grid. The program is shown in Fig. 8. Processes `P1` and `P2` should always access disjoint memory locations and hence no synchronization between the processes should be required. `P1` starts by reading a value from the input and then begins a loop which adds this value to the locations on the diagonal of the 2D matrix. We iterate over the diagonal of the 2D (x,y) plane as the value of variable `z1` is fixed to 1 and only `x1` and `y1` change. The loop comprises the statements at labels L2 and L6 in `P1`. After the plane is updated, `P1` updates a value in another plane (L7-L9). For clarity we have only shown the update of a single location, but this can also be extended to update the diagonal. Similarly, process `P2` updates the diagonal

```
Barrier = F1 = F2 = 0;
Region = 2;
Defragment() {
 /* Pick a Region */
 L1: i = Region;
 L2: Region = i - 1;
 L3: empty = i - 2;
 L4: if (i >= N) goto L14;
 /* has free entry? */
 L5: b = Pages[i];
 L6: if (!b  && empty <= 0)
 L7:  empty = i;
 /* Copy Entry */
 L8: if (b  && empty > 0) {
 L9:  Pages[empty] = true;
 L10: empty += 2;
 L11: Pages[i] = false;
   }
 L12: i += 2;
 L13: goto L4;
 /* Barrier Synch */
 L14: Barrier += 1; F1 = 0;
 L15: if (F1 == 1)
    goto L16;
    if (Barrier == 2) {
    Barrier = 0; F2 = 1;
    Region = 2;
    goto L16;
    }
    goto L15;
 L16: goto L1;
}

Update() {
 /* Pick a Region */
 L1: j = Region;
 L2: Region = j - 1;
 L3: b = Pages[j];
 /* Update the Page */
 L4: if (!b)
    Pages[j] = true;
 /* Barrier Sync */
 L5: Barrier += 1; F2 = 0;
 L6: if (F2 == 1)
    goto L7;
   if (Barrier == 2) {
    Barrier = 0; F1 = 1;
    Region = 2;
    goto L7;
   }
   goto L6;
 L7:
}

main() {
  Defragment() || Update();
}
```

Figure 7. Concurrent Defragmentation

```
x1 = 0; y1 = 0; z1 = 1;
x2 = 0; y2 = 1; z2 = 1;
x3 = 0; y3 = 1; z3 = 0;

P1()
 L1: v = read();
 L2: r = A[x1][y1][z1];
 L3: A[x1][y1][z1] = r + v;
 L4: x1 += 1;
 L5: y1 += 1;
 L6: if (x1 < N)
    goto L2;

 L7: v = read();
 L8: r = A[x3][y3][z3];
 L9: A[x3][y3][z3] = r + v;

P2()
 L1: v = read();
 L2: r = A[x2][y2][z2];
 L3: A[x2][y2][z2] = r + v;
 L4: y2 += 1;
 L5: z2 += 1;
 L6: if (y2 < N)
    goto L2;

main()
 P1() || P2()
```

Figure 8. Concurrent 3D Updating

of a 2D plane but this time in the (z,y) dimension. That is, the value of `x2` is fixed and only `y2` and `z2` are updated.

Specification: The two processes should never access the same locations simultaneously. That is, if process `P1` is reading or writing shared data (e.g. at labels L2, L3, L8, L9), `P2` should not be writing simultaneously (e.g. be at L3), and vice versa for `P2`, where the indices of the array being accessed are equal for both processes.

Result: As shown in Table 2, refining the abstraction leads to weaker atomicity constraints. In this example, we have 3 layers of abstractions, each leading to finer-grained solutions.

8. Related Work

Synthesis from Temporal Specifications: Early work by Emerson and Clarke [7] uses temporal specifications to generate a synchronization skeleton. This has been extended by Attie and Emerson to synthesize programs with finer grained atomic sections [2]. Early work by Manna and Wolper [16] synthesizes CSP programs. Pnueli and Rosner [20] consider the problem of synthesizing a reactive module based on an LTL specification. They discuss the problem of *implementability* in this setting, and define necessary and sufficient conditions for the implementability of a given specification. Our work focuses on concurrent programs for shared memory and is based on abstract interpretation, handling infinite-state systems.

Program Repair and Game-Based Synthesis: Jobstmann et. al. [14] consider the problem of *program repair* as a game. In their approach, a game is constructed from (a modified version of) the program to be repaired, and an LTL specification of the correctness property. The problem of repair boils down to finding a winning strategy in that game. This approach has been later extended to provide fault localization and fixing [27, 15]. The approach has also been extended to work with predicate abstraction in [12]. In contrast to these, we focus on concurrent programs, use abstract interpretation, and solve the quantitative problem of computing a minimally constrained program.

In our previous work [30], we focused on inference of CCR guards in finite-state concurrent programs, where the atomic blocks were not modified. This work can be viewed as the next general step and addresses the more general problem of infinite-state systems, employs abstract interpretation, and infers atomicity constraints (as opposed to only inferring guards).

Dynamic Approaches: The problem of restricting the program to valid executions can be addressed by monitoring the program at *runtime* and forcing it to avoid executions that violate the specification. However, restricting the executions of a program at runtime requires a recovery mechanism in case the program already performed a step that violates the specification, and/or a predictive mechanism to check whether future steps lead to a violation.

Existing approaches using recovery mechanisms typically require user annotations to define a corrective action to be taken when the specification is violated. For example, software transactional memory [23] is a special case of a recovery mechanism in which the user provides atomicity annotations defining atomic sections. The system then requires the absence of read/write conflicts, and if this property is violated, the execution of an atomic section is restarted. Other examples include Tolerace [19] which creates local copies of variables to detect and recover from races, and ISOLATOR [21] which can recover from violations of isolation.

Search-based Synthesis: In previous work [29, 28], we used a semi-automated approach for exploring a space concurrent garbage collectors and linearizable data-structures. The work used a search procedure and an abstraction specifically geared towards the safety property required for the specific domain.

In *sketching* [26, 25], the user provides a reference program of the desired implementation and some sketches which partially specify certain optimized functions. The sketching compiler automatically fills in the missing low-level details to create an optimized implementation. Sketching has been used for bounded programs and in special cases of unbounded domains [24]. In [25], finding a candidate solution is done using a counterexample-guided inductive synthesis (CEGIS) algorithm that uses a backing bounded-checking procedure. Candidates are generated iteratively and run through the checker. Counterexamples are used to limit the next candidates to be generated. In contrast, rather than generating candidates and checking them, in our approach, the synthesizer is part of the verification procedure and is based on abstract interpretation. Further, in contrast to sketching, which aims to find *some* solution for the sketch, we are interested in finding a solution with minimal synchronization.

Locks for Atomicity: There have been several works on inferring locks for atomic sections. In the work by McCloskey et. al. [17], a tool called Autolocker is presented. The tool takes as input a pro-

gram that has been manually annotated with (i) atomic sections and (ii) a mapping between locks and memory locations protected by these locks. Autolocker produces a program that implements the atomic sections in (i) with the locks in (ii). Further work by Emmi et. al. [11] proposed a technique to automate part (ii) above. The actual assignment of locations to locks is solved as an optimization problem where the goal is to minimize the total number of locks while still achieving minimum interference between the computed locks. The latest work of Cherem et. al. [6] proposes another alternative to automate (ii) while also computing actual lock placement in the code. Our work is complementary to these approaches, as our focus is not on optimizing the implementation of atomic sections, but on inferring minimally atomic synchronization.

9. Conclusions and Future Work

In this paper, we presented a novel algorithm for the automatic synthesis of efficient synchronization in concurrent infinite-state programs (AGS). The synchronization can be realized by modifying either the program or the scheduler. Our algorithm is based on abstract interpretation and thus applies to concurrent infinite-state programs.

The AGS algorithm leads to a new verification approach: it allows for both the abstraction *and* the program to be modified simultaneously until the abstraction is precise enough to verify the (modified) program. This enables verification of a program in cases where it would have otherwise failed.

We implemented the AGS approach in a prototype tool named GUARDIAN, and successfully applied it to several small but interesting concurrent programs. GUARDIAN works with various numerical abstractions. In the future, we intend to investigate its use with finer abstract domains, such as the trace partitioning domain [22], which is a natural fit for our setting, as it allows to abstract interleavings with varying degrees of precision.

We demonstrated our approach using atomic sections as the synchronization primitive, but `avoid` and `implement` can be realized using other synchronization primitives. In the future, we intend to explore extensions of AGS to other synchronization primitives.

The AGS algorithm described in this paper can also be applied in a dynamic setting, where invalid interleavings are obtained by running the program driven by test-cases. In such a setting, the constraints obtained from dynamic executions can be used to give the user partial program corrections, or used to limit the space that has to be explored statically.

Acknowledgement

The authors wish to thank Mooly Sagiv for many insightful comments on an earlier version of this work.

References

[1] The SAT4J SAT solver. available at `http://www.sat4j.org/`.

[2] ATTIE, P., AND EMERSON, E. Synthesis of concurrent systems for an atomic read/atomic write model of computation. In *PODC '96* (1996), ACM, pp. 111–120.

[3] BALL, T., AND RAJAMANI, S. K. Automatically validating temporal safety properties of interfaces. In *SPIN* (2001), pp. 103–122.

[4] BLANCHET, B., COUSOT, P., COUSOT, R., FERET, J., MAUBORGNE, L., MINÉ, A., MONNIAUX, D., AND RIVAL, X. A static analyzer for large safety-critical software. In *PLDI* (2003), pp. 196–207.

[5] BLOEM, R., CHATTERJEE, K., HENZINGER, T., AND JOBSTMANN, B. Better quality in synthesis through quantitative objectives. In *CAV* (2009), pp. 140–156.

[6] CHEREM, S., CHILIMBI, T., AND GULWANI, S. Inferring locks for atomic sections. In *PLDI* (2008), pp. 304–315.

[7] CLARKE, E., AND EMERSON, E. Design and synthesis of synchronization skeletons using branching-time temporal logic. In *Logic of Programs, Workshop* (1982), pp. 52–71.

[8] CLARKE, E. M., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. Counterexample-guided abstraction refinement. In *CAV* (2000), pp. 154–169.

[9] CLARKE, JR., E., GRUMBERG, O., AND PELED, D. *Model Checking*. The MIT Press, 1999.

[10] COUSOT, P., AND COUSOT, R. Abstract interpretation: A unified lattice model for static analysis of programs by construction of approximation of fixed points. In *POPL* (1977), pp. 238–252.

[11] EMMI, M., FISCHER, J. S., JHALA, R., AND MAJUMDAR, R. Lock allocation. In *POPL* (2007), pp. 291–296.

[12] GRIESMAYER, A., BLOEM, R. P., AND COOK, B. Repair of boolean programs with an application to C. In *CAV* (2006), pp. 358–371.

[13] HENZINGER, T. A., JHALA, R., MAJUMDAR, R., AND SUTRE, G. Lazy abstraction. In *POPL* (2002), pp. 58–70.

[14] JOBSTMANN, B., GRIESMAYER, A., AND BLOEM, R. Program repair as a game. In *CAV* (2005), pp. 226–238.

[15] JOBSTMANN, B., STABER, S., GRIESMAYER, A., AND BLOEM, R. Finding and fixing faults. *Journal of Computer and System Sciences (JCSS)* (2008).

[16] MANNA, Z., AND WOLPER, P. Synthesis of communicating processes from temporal logic specifications. *ACM Trans. Program. Lang. Syst. (TOPLAS) 6*, 1 (1984), 68–93.

[17] MCCLOSKEY, B., ZHOU, F., GAY, D., AND BREWER, E. Autolocker: synchronization inference for atomic sections. In *POPL* (2006), pp. 346–358.

[18] MINÉ, A. The octagon abstract domain. *Higher Order Symbol. Comput. 19*, 1 (2006), 31–100.

[19] NAGPALY, R., PATTABIRAMANZ, K., KIROVSKI, D., AND ZORN, B. Tolerace: Tolerating and detecting races. In *STMCS: Second Workshop on Software Tools for Multi-Core Systems* (2007).

[20] PNUELI, A., AND ROSNER, R. On the synthesis of a reactive module. In *POPL '89* (New York, NY, USA, 1989), ACM, pp. 179–190.

[21] RAJAMANI, S., RAMALINGAM, G., RANGANATH, V.-P., AND VASWANI, K. Controlling non-determinism for semantic guarantees. In *Exploiting Concurrency Efficiently and Correctly – (EC)2* (2008).

[22] RIVAL, X., AND MAUBORGNE, L. The trace partitioning abstract domain. *ACM Trans. Program. Lang. Syst. 29*, 5 (2007), 26.

[23] SHAVIT, N., AND TOUITOU, D. Software transactional memory. In *PODC '95* (New York, NY, USA, 1995), ACM, pp. 204–213.

[24] SOLAR-LEZAMA, A., ARNOLD, G., TANCAU, L., BODÍK, R., SARASWAT, V. A., AND SESHIA, S. A. Sketching stencils. In *PLDI* (2007), pp. 167–178.

[25] SOLAR-LEZAMA, A., JONES, C. G., AND BODIK, R. Sketching concurrent data structures. In *PLDI* (2008), pp. 136–148.

[26] SOLAR-LEZAMA, A., RABBAH, R. M., BODÍK, R., AND EBCIOGLU, K. Programming by Sketching for Bit-Streaming Programs. In *PLDI* (2005), pp. 281–294.

[27] STABER, S., JOBSTMANN, B., AND BLOEM, R. Finding and fixing faults. In *CHARME* (2005), pp. 35–49.

[28] VECHEV, M., AND YAHAV, E. Deriving linearizable fine-grained concurrent objects. In *PLDI* (2008), pp. 125–135.

[29] VECHEV, M. T., YAHAV, E., BACON, D. F., AND RINETZKY, N. Cgcexplorer: a semi-automated search procedure for provably correct concurrent collectors. In *PLDI* (2007), pp. 456–467.

[30] VECHEV, M. T., YAHAV, E., AND YORSH, G. Inferring synchronization under limited observability. In *TACAS* (2009), pp. 139–154.