# Automated Synthesis of Asynchronizations

Sidi Mohamed Beillahi[1], Ahmed Bouajjani[2], Constantin Enea[3], and Shuvendu Lahiri[4]

[1] University of Toronto, Canada
sm.beillahi@utoronto.ca
[2] Université Paris Cité, IRIF, CNRS, Paris, France
abou@irif.fr
[3] LIX, Ecole Polytechnique, CNRS and Institut Polytechnique de Paris, France
cenea@lix.polytechnique.fr
[4] Microsoft Research Lab - Redmond
shuvendu@microsoft.com

**Abstract.** Asynchronous programming is widely adopted for building responsive and efficient software, and modern languages such as C# provide async/await primitives to simplify the use of asynchrony. In this paper, we propose an approach for refactoring a sequential program into an asynchronous program that uses async/await, called asynchronization. The refactoring process is parametrized by a set of methods to replace with asynchronous versions, and it is constrained to avoid introducing data races. We investigate the delay complexity of enumerating all data race free asynchronizations, which quantifies the delay between outputting two consecutive solutions. We show that this is polynomial time modulo an oracle for solving reachability in sequential programs. We also describe a pragmatic approach based on an interprocedural dataflow analysis with polynomial-time delay complexity. The latter approach has been implemented and evaluated on a number of non-trivial C# programs extracted from open-source repositories.

## 1 Introduction

Asynchronous programming is widely adopted for building responsive and efficient software. As an alternative to explicitly registering callbacks with asynchronous calls, C# 5.0 [3] introduced the `async`/`await` primitives. These primitives allow the programmer to write code in a familiar sequential style without explicit callbacks. An asynchronous procedure, marked with `async`, returns a task object that the caller uses to "await" it. Awaiting may suspend the execution of the caller, but does not block the thread it is running on. The code after `await` is the continuation called back when the callee result is ready. This paradigm has become popular across many languages, C++, JavaScript, Python.

The `async`/`await` primitives introduce concurrency which is notoriously complex. The code in between a call and a matching `await` (referring to the same task) may execute before some part of the awaited task or after the awaited task finished. For instance, on the middle of Fig. 1, the assignment `y=1` at line 4 can execute before or after `RdFile` finishes. The `await` for `ReadToEndAsync` in

```
1 void Main(string f) {
2  x = 0;
3  int val = RdFile(f);
4  y = 1;

6  int r = x;
7  Debug.Assert(r == val); }

9 int RdFile(string f) {
10 var rd=new StreamReader(f);
11 string s = rd.ReadToEnd();
12 int r1 = x;

14 x = r1 + s.Length;
15 return s.Length;    }
```

```
1 async Task Main(string f) {
2  x = 0;
3  Task<int> t1 = RdFile(f);
4  y = 1;
5  int val = await t1;
6  int r = x;
7  Debug.Assert(r == val);  }

9 async Task<int> RdFile(string f) {
10 var rd = new StreamReader(f);
11 Task<string> t=rd.ReadToEndAsync();
12 int r1 = x;
13 string s = await t;
14 x = r1 + s.Length;
15 return s.Length;              }
```

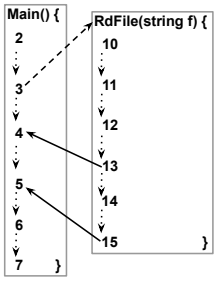


Fig. 1: Synchronous and asynchronous C# programs (`x`, `y` are static variables).

`RdFile` (line 13) may suspend `RdFile`'s execution because `ReadToEndAsync` did not finish, and pass the control to `Main` which executes `y=1`. If `ReadToEndAsync` finishes before this `await` executes, then the latter has no effect and `y=1` gets executed after `RdFile` finishes. The resemblance with sequential code can be especially deceitful since this non-determinism is opaque. It is common that `await`s are placed immediately after the corresponding call which limits the benefits that can be obtained from executing steps in the caller and callee concurrently [24].

In this paper, we address the problem of writing efficient asynchronous code that uses `async`/`await`. We propose a procedure for automated synthesis of asynchronous programs *equivalent* to a given synchronous (sequential) program $P$. This can be seen as a way of refactoring synchronous code to asynchronous code. Solving this problem in its full generality would require checking equivalence between arbitrary programs, which is known to be hard. Therefore, we consider a restricted space of asynchronous program candidates defined by substituting synchronous methods in $P$ with asynchronous versions (assumed to be behaviorally equivalent). The substituted methods are assumed to be leaves of the call-tree (they do not call any method in $P$). Such programs are called *asynchronizations* of $P$. A practical instantiation is replacing IO synchronous calls for reading/writing files or managing http connections with asynchronous versions.

For instance, the sequential C# program on the left of Fig. 1 contains a `Main` that invokes a method `RdFile` that returns the length of the text in a file. The file name input to `RdFile` is an input to `Main`. The program uses a variable `x` to aggregate the lengths of all files accessed by `RdFile`; this would be more useful when `Main` calls `RdFile` multiple times which we omit for simplicity. Note that this program passes the assertion at line 7. The time consuming method `ReadToEnd` for reading a file is an obvious choice for being replaced with an equivalent *asynchronous* version whose name is suffixed with `Async`. Performing such tasks asynchronously can lead to significant performance boosts. The program on the middle of Fig. 1 is an example of an asynchronization defined by this substitution. The syntax of `async`/`await` imposes that every method that transitively calls one of the substituted methods, i.e., `Main` and `RdFile`, must also be declared as asynchronous. Then, every asynchronous call must be followed by an

`await` that specifies the control location where that task should have completed. For instance, the `await` for `ReadToEndAsync` is placed at line 13 since the next instruction (at line 14) uses the computed value. Therefore, synthesizing such refactoring reduces to finding a correct placement of `await`s (that implies equivalence) for every call of a method that transitively calls a substituted method (we do not consider "deeper" refactoring like rewriting conditionals or loops).

We consider an equivalence relation between a synchronous program and an asynchronization that corresponds to absence of data races in the asynchronization. Data race free asynchronizations are called *sound*. Relying on absence of data races avoids reasoning about equality of sets of reachable states which is harder in general, and an established compromise in reasoning about concurrency. For instance, the asynchronization in Fig. 1 is sound because the call to `RdFile` accessing `x` finishes before the read of `x` in `Main` (line 6). Therefore, accesses to `x` are performed in the same order as in the synchronous program.

The asynchronization on the right of Fig. 1 is not the only sound (data-race free) asynchronization of the program on the left. The `await` at line 13 can be moved one statement up (before the read of `x`) and the resulting program remains equivalent to the sequential one. In this paper, we investigate the problem of enumerating *all* sound asynchronizations of a sequential program $P$ w.r.t. substituting a set of methods with asynchronous versions. This makes it possible to deal separately with the problem of choosing the best asynchronization in terms of performance based on some metric (e.g., performance tests).

Identifying the most efficient asynchronization is difficult and can not be done syntactically. It is tempting to consider that increasing the distance between calls and matching `await`s so that more of the caller code is executed while waiting for an asynchronous task to finish increases performance. However, this is not true in general. We use the programs in Fig. 2 to show that the best `await` placement w.r.t. performance depends on execution times of code blocks in between calls and `await`s in a non-trivial manner. Note that estimating these execution times, especially for IO operations like http connections, can not be done statically.

The programs in Fig. 2 use `Thread.Sleep(n)` to abstract sequential code executing in $n$ milliseconds and `Task.Delay(n)` to abstract an asynchronous call executing in $n$ milliseconds on a different thread. The functions named `Foo` differ only in the position of `await t`. We show that modifying this position worsens execution time in each case. For the left program, best performance corresponds to maximal distance between `await t` in `Foo` and the corresponding call. This allows the `IO` call to execute in parallel with the caller, as depicted on the bottom-left of Fig. 2. The executions corresponding to the other two positions of `await t` are given just above. For the middle program, placing `await t` in between the two code blocks in `Foo` optimizes performance (note the extra `IO` call in `Main`): the `IO` call in `Foo` executes in parallel with the first code block in `Foo` and the `IO` call in `Main` executes in parallel with the second one. This is depicted on the bottom-middle of Fig. 2. The execution above shows that placing `await t` as on the left (after the two code blocks) leads to worse execution time (placing `await t` immediately after the call is also worse). Finally, for the right program, placing

```
1 async Task Main() {
2  var t1 = Foo();

4  await t1;
5                  }

7 async Task Foo() {
8  var t = IO();

10 Thread.Sleep(200);

12 Thread.Sleep(200);
13 await t;        }

15 async Task IO() {
16  var t0 = Task.Delay(300);

18  await t0;      }
```

```
1 async Task Main() {
2  var t1 = Foo();
3  var t2 = IO();
4  await t1;
5  await t2;        }

7 async Task Foo() {
8  var t = IO();

10 Thread.Sleep(200);
11 await t;
12 Thread.Sleep(200);
13                 }

15 async Task IO() {
16  var t0 = Task.Delay(300);

18  await t0;      }
```

```
1 async Task Main() {
2  var t1 = Foo();
3  var t2 = IO();
4  await t1;
5  await t2;        }

7 async Task Foo() {
8  var t = IO();
9  await t;
10 Thread.Sleep(200);

12 Thread.Sleep(200);
13                 }

15 async Task IO() {
16  var t0 = Task.Delay(300);
17  Thread.Sleep(150);
18  await t0;      }
```

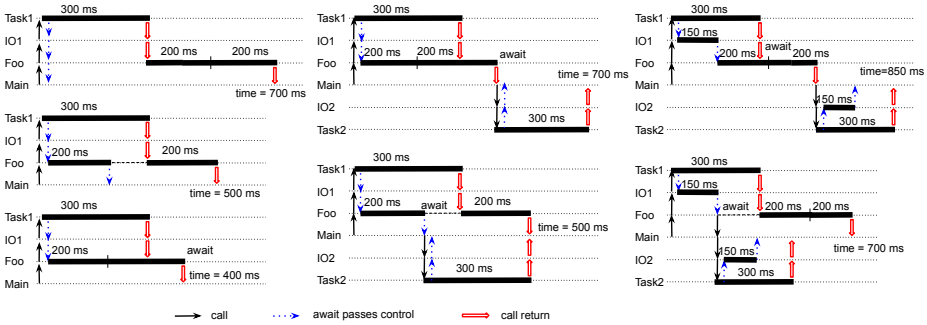


Fig. 2: Asynchronous C# programs and executions. On the bottom, time durations of executing code blocks from the same method are aligned horizontally, and time goes from left to right. Vertical single-line arrows represent method call steps, dashed arrows represent `await`s passing control to the caller, and double-line arrows represent a call return. Total execution time is marked `time=...`.

`await t` immediately after the call is best (note that `IO` executes another code block before `await`). The `IO` call in `Main` executes in parallel with `Foo` as shown on the bottom-right of Fig. 2. The execution above shows the case where `await t` is placed in the middle (the await has no effect because `IO` already finished, and `Foo` continues to execute). This leads to worse execution time (placing `await t` after the two code blocks is also worse). These differences in execution times have been confirmed by running the programs on a real machine.

As demonstrated by the examples in Fig. 2, the performance of an asynchronization depends on the execution environment, e.g., the overhead of IO operations like http connections and disk access (in Fig. 2, we use `Thread.Sleep(n)` or `Task.Delay(n)` to model such overheads). Since modeling the behavior of an execution environment w.r.t. performance is difficult in general, selecting the most performant asynchronization using static reasoning is also difficult. As a way of sidestepping this difficulty, we focus on enumerating *all* sound asynchronizations that allows to evaluate performance separately in a dynamic manner using performance tests for instance (for each sound asynchronization).

In the worst-case, the number of (sound) asynchronizations is exponential in the number of method calls in the program. Therefore, we focus on the *delay complexity* of the problem of enumerating sound asynchronizations, i.e., the complexity of the delay between outputting two consecutive (distinct) solutions, and show that this is polynomial time modulo an oracle for solving reachability (assertion checking) in *sequential* programs. Note that a trivial enumeration of all asynchronizations and checking equivalence for each one of them has an exponential delay complexity modulo an oracle for checking equivalence.

As an intermediate step, we consider the problem of computing *maximal* sound asynchronizations that maximize the distance between every call and its matching `await`. We show that rather surprisingly, there exists a *unique* maximal sound asynchronization. This is not trivial since asynchronizations can be incomparable w.r.t. distances between calls and `await`s (i.e., better for one `await` and worse for another, and vice-versa). This holds even if maximality is relative to a given asynchronization $P_a$ imposing an upper bound on the distance between awaits and calls. In principle, avoiding data races could reduce to a choice between moving one await or another closer to the matching call. We show that this is not necessary because the maximal asynchronization is required to be equivalent to a *sequential* program, which executes statements in a fixed order.

As a more pragmatic approach, we define a procedure for computing sound asynchronizations which relies on a bottom-up interprocedural data-flow analysis. The placement of awaits is computed by traversing the call graph bottom up and using a data-flow analysis that computes read or write accesses made in the callees. We show that this procedure computes maximal sound asynchronizations of abstracted programs where every Boolean condition is replaced with a non-deterministic choice. These asynchronizations are sound for the concrete programs as well. This procedure enables a polynomial-time delay enumeration of sound asynchronizations of abstracted programs.

We implemented the asynchronization enumeration based on data-flow analysis in a prototype tool for C# programs. We evaluated this implementation on a number of non-trivial programs extracted from open source repositories to show that our techniques have the potential to become the basis of refactoring tools that allow programmers to improve their usage of async/await primitives.

In summary, this paper makes the following contributions:

- Define the problem of data race-free (sound) asynchronization synthesis for refactoring sequential code to equivalent asynchronous code (Section 3).
- Show that the problem of computing a sound asynchronization that maximizes the distance between calls and awaits has a unique solution (Section 4).
- The delay complexity of sound asynchronization synthesis (Sections 5–6).
- A pragmatic algorithm for computing sound asynchronizations based on a data-flow analysis (Section 7).
- A prototype implementation of this algorithm and an evaluation of this prototype on a benchmark of non-trivial C# programs (Section 8).

Additional formalization and proofs are included in the appendix.

$$
\begin{array}{lll}
\langle prog \rangle & ::= & \texttt{program}\ \langle md \rangle \\
\langle md \rangle & ::= & \texttt{method}\ \langle m \rangle\ \{\ \langle inst \rangle\ \} \mid \texttt{async method}\ \langle m \rangle\ \{\ \langle inst \rangle\ \} \mid \langle md \rangle\ \langle md \rangle \\
\langle inst \rangle & ::= & \langle x \rangle\ \texttt{:=}\ \langle le \rangle \mid \langle r \rangle\ \texttt{:=}\ \langle x \rangle \mid \langle r \rangle := \texttt{call}\ \langle m \rangle \mid \texttt{return} \mid \texttt{await}\ \langle r \rangle \\
 & & \mid \texttt{await}\ * \mid \texttt{if}\ \langle le \rangle\ \{\langle inst \rangle\}\ \texttt{else}\ \{\langle inst \rangle\} \mid \texttt{while}\ \langle le \rangle\ \{\langle inst \rangle\} \mid \\
 & & \langle inst \rangle\ ;\ \langle inst \rangle
\end{array}
$$

Fig. 3: Syntax. $\langle m \rangle$, $\langle x \rangle$, and $\langle r \rangle$ represent method names, program and local variables, resp. $\langle le \rangle$ is an expression over local variables, or $*$ which is non-deterministic choice.

## 2 Asynchronous Programs

We consider a simple programming language to formalize our approach, shown in Fig. 3. A *program* is a set of methods, including a distinguished `main`, which are classified as *synchronous* or *asynchronous*. Synchronous methods run continuously until completion when they are invoked. Asynchronous methods, marked using the keyword `async`, can run only partially and be interrupted when executing an `await`. Only asynchronous methods can use `await`, and all methods using `await` must be defined as asynchronous. We assume that methods are not (mutually) recursive. A program is called *synchronous* if it is a set of synchronous methods.

A method is defined by a name from a set $\mathbb{M}$ and a list of statements over a set $\mathbb{PV}$ of *program variables*, which can be accessed from different methods (ranged over using $x$, $y$, $z$,...), and a set $\mathbb{LV}$ of method *local variables* (ranged over using $r$, $r_1$, $r_2$,...). Input/return parameters are modeled using program variables. Each method call returns a *unique task identifier* from a set $\mathbb{T}$, used to record control dependencies imposed by `await`s (for uniformity, synchronous methods return a task identifier as well). Our language includes assignments, `await`s, `return`s, loops, and conditionals. Assignments to a local variable $r := x$, where $x$ is a program variable, are called *reads* of $x$, and assignments to a program variable $x := le$ ($le$ is an expression over local variables) are called *writes* to $x$. A *base* method is a method whose body does *not* contain method calls.

**Asynchronous methods.** Asynchronous methods can use `await`s to wait for the completion of a task (invocation) while *the control is passed to their caller*. The parameter $r$ of the `await` specifies the id of the awaited task. As a sound abstraction of awaiting the completion of an IO operation (reading or writing a file, an http request, etc.), which we do not model explicitly, we use a variation `await` $*$. This has a non-deterministic effect of either continuing to the next statement in the same method (as if the IO operation already completed), or passing the control to the caller (as if the IO operation is still pending).

```
async method ReadToEndAsync() {
 await *;
 ind = Stream.index;
 len = Stream.content.Length;
 if (ind >= len)
  retVal  =  ""; return
 Stream.index = len;
 retVal = Stream.content(ind,len);
 return                          }
```

Fig. 4: An IO method.

Fig. 4 lists our modeling of the IO method `ReadToEndAsync` used in Fig. 1. We use program variables to represent system resources such as the file system. The await for the completion of accesses to such resources is modeled by `await`

$*$. This enables capturing racing accesses to system resources in asynchronous executions. Parameters or return values are modeled using program variables. `ReadToEndAsync` is modeled using reads/writes of the index/content of the input stream, and `await` $*$ models the await for their completion.

We assume that the body of every asynchronous method $m$ satisfies several well-formedness syntactic constraints, defined on its control-flow graph (CFG). We recall that each node of the CFG represents a basic block of code (a maximal-length sequence of branch-free code), and nodes are connected by directed edges which represent a possible transfer of control between blocks. Thus,

1. every call $r :=$ `call` $m'$ uses a distinct variable $r$ (to store task identifiers),
2. every CFG block containing an `await` $r$ is dominated by the CFG block containing the call $r :=$ `call` ... (i.e., every CFG path from the entry to the await has to pass through the call),
3. every CFG path starting from a block containing a call $r :=$ `call` ... to the exit has to pass through an `await` $r$ statement.

The first condition simplifies the technical exposition, while the last two ensure that $r$ stores a valid task identifier when executing an `await` $r$, and that every asynchronous invocation is awaited before the caller finishes. Languages like C# or Javascript do not enforce the latter constraint, but it is considered bad practice due to possible exceptions that may arise in the invoked task and are not caught. We forbid passing task identifiers as method parameters (which is possible in C#). A statement `await` $r$ is said to *match* a statement $r :=$ `call` $m'$.

In Fig. 5, we give three examples of programs to explain in more details the well-formedness syntactic constraints. The program on the left of Fig. 5 does not satisfy the second condition since `await r` can be reached without entering the loop. The program in the center of Fig. 5 does not satisfy the third condition since we can reach the end of the method without entering the if branch and thus, without executing `await r`. The program on the right of Fig. 5 satisfies both conditions.

```
async method m {
  while *
    r = call m1;

  await r;
}
```

```
async method m {
  r = call m1;

  if *
    await r;
}
```

```
async method m {
  r = call m1;
  while *
    r' = call m1;
    await r';
  await r;
}
```

Fig. 5: Examples of programs

**Semantics.** A program configuration is a tuple $(\mathsf{g}, \mathsf{stack}, \mathsf{pend}, \mathsf{cmpl}, \mathsf{c\text{-}by}, \mathsf{w\text{-}for})$ where $\mathsf{g}$ is composed of the valuation of the program variables excluding the program counter, $\mathsf{stack}$ is the call stack, $\mathsf{pend}$ is the set of asynchronous tasks, e.g., continuations predicated on the completion of some method call, $\mathsf{cmpl}$ is the set of completed tasks, $\mathsf{c\text{-}by}$ represents the relation between a method call and its caller, and $\mathsf{w\text{-}for}$ represents the control dependencies imposed by `await` statements. The activation frames in the call stack and the asynchronous tasks are represented using triples $(i, m, \ell)$ where $i \in \mathbb{T}$ is a task identifier, $m \in \mathbb{M}$ is a method name, and $\ell$ is a valuation of local variables, including as usual a dedicated program counter. The set of completed tasks is represented as a function $\mathsf{cmpl} : \mathbb{T} \to \{\top, \bot\}$ such that $\mathsf{cmpl}(i) = \top$ when $i$ is completed and $\mathsf{cmpl}(i) = \bot$, otherwise. We define $\mathsf{c\text{-}by}$ and $\mathsf{w\text{-}for}$ as partial functions $\mathbb{T} \rightharpoonup \mathbb{T}$ with

the meaning that $\mathsf{c\text{-}by}(i) = j$, resp., $\mathsf{w\text{-}for}(i) = j$, iff $i$ is called by $j$, resp., $i$ is waiting for $j$. We set $\mathsf{w\text{-}for}(i) = *$ if the task $i$ was interrupted because of an `await` $*$ statement.

The semantics of a program $P$ is defined as a labeled transition system (LTS) $[P] = (\mathbb{C}, \mathbb{A}\mathsf{ct}, \mathsf{ps}_0, \rightarrow)$ where $\mathbb{C}$ is the set of program configurations, $\mathbb{A}\mathsf{ct}$ is a set of transition labels called *actions*, $\mathsf{ps}_0$ is the initial configuration, and $\rightarrow \subseteq \mathbb{C} \times \mathbb{A}\mathsf{ct} \times \mathbb{C}$ is the transition relation. Each program statement is interpreted as a transition in $[P]$. The set of actions is defined by ($\mathbb{A}\mathsf{id}$ is a set of action identifiers):

$$\begin{aligned}\mathbb{A}\mathsf{ct} = \{&(aid, i, ev) : aid \in \mathbb{A}\mathsf{id}, i \in \mathbb{T}, ev \in \{\mathsf{rd}(x), \mathsf{wr}(x), \mathsf{call}(j), \mathsf{await}(k), \mathsf{return},\\ &\mathsf{cont} : j \in \mathbb{T}, k \in \mathbb{T} \cup \{*\}, x \in \mathbb{PV}\}\}\end{aligned}$$

The transition relation $\rightarrow$ is defined in Fig. 6. Transition labels are written on top of $\rightarrow$.

Transitions labeled by $(aid, i, \mathsf{rd}(x))$ and $(aid, i, \mathsf{wr}(x))$ represent a read and a write accesses to the program variable $x$, respectively, executed by the task (method call) with identifier $i$. A transition labeled by $(aid, i, \mathsf{call}(j))$ corresponds to the fact that task $i$ executes a method call that results in creating a task $j$. Task $j$ is added on the top of the stack of currently executing tasks, declared pending (setting $\mathsf{cmpl}(j)$ to $\bot$), and $\mathsf{c\text{-}by}$ is updated to track its caller ($\mathsf{c\text{-}by}(j) = i$). A transition $(aid, i, \mathsf{return})$ represents the return from task $i$. Task $i$ is removed from the stack of currently executing tasks, and $\mathsf{cmpl}(i)$ is set to $\top$ to record the fact that task $i$ is finished.

A transition $(aid, i, \mathsf{await}(j))$ relates to task $i$ waiting asynchronously for task $j$. Its effect depends on whether task $j$ is already completed. If this is the case (i.e., $\mathsf{cmpl}[j] = \top$), task $i$ continues and executes the next statement. Otherwise, task $i$ executing the $\mathsf{await}$ is removed from the stack and added to the set of pending tasks, and $\mathsf{w\text{-}for}$ is updated to track the waiting-for relationship ($\mathsf{w\text{-}for}(i) = j$). Similarly, a transition $(aid, i, \mathsf{await}(*))$ corresponds to task $i$ waiting asynchronously for the completion of an unspecified task. Nondeterministically, task $i$ continues to the next statement, or task $i$ is interrupted and transferred to the set of pending tasks ($\mathsf{w\text{-}for}(i)$ is set to $*$).

A transition $(aid, i, \mathsf{cont})$ represents the scheduling of the continuation of task $i$. There are two cases depending on whether $i$ waited for the completion of another task $j$ modeled explicitly in the language (i.e., $\mathsf{w\text{-}for}(i) = j$), or an unspecified task (i.e., $\mathsf{w\text{-}for}(i) = *$). In the first case, the transition is enabled only when the call stack is empty and $j$ is completed. In the second case, the transition is always enabled. The latter models the fact that methods implementing IO operations (waiting for unspecified tasks in our language) are executed in background threads and can interleave with the main thread (that executes the `Main` method). Although this may seem restricted because we do not allow arbitrary interleavings between IO methods and `Main`, this is actually sound when focusing on the existence of data races as in our approach. As shown later in Table 1, any two instructions that follow an `await` $*$ are not happens-before related and form a race.

$$\frac{\mathtt{r} := \mathtt{x} \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \ell' = \ell[r \mapsto \mathsf{g}(x), \mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \_, \_, \_, \_) \xrightarrow{(aid,\, i,\, \mathsf{rd}(x))} (\mathsf{g}, (i, m, \ell') \circ \mathsf{stack}, \_, \_, \_, \_)}$$

$$\frac{\mathtt{x} := \mathtt{1e} \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \ell' = \ell[\mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))] \quad \mathsf{g}' = \mathsf{g}[x \mapsto \ell(\mathtt{1e})]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \_, \_, \_, \_) \xrightarrow{(aid,\, i,\, \mathsf{wr}(x))} (\mathsf{g}', (i, m, \ell') \circ \mathsf{stack}, \_, \_, \_, \_)}$$

$$\frac{\begin{array}{c} r := \mathsf{call}\ m \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \ell_0 = \mathsf{init}(\mathsf{g}, m) \quad j \in \mathbb{T}\ \text{fresh} \\ \ell' = \ell[r \mapsto j, \mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))] \quad \mathsf{cmpl}' = \mathsf{cmpl}[j \mapsto \bot] \quad \mathsf{c\text{-}by}' = \mathsf{c\text{-}by}[j \mapsto i] \end{array}}{(\mathsf{g}, (i, m', \ell) \circ \mathsf{stack}, \_, \mathsf{cmpl}, \mathsf{c\text{-}by}, \_) \xrightarrow{(aid,\, i,\, \mathsf{call}(j))} (\mathsf{g}, (j, m, \ell_0) \circ (i, m', \ell') \circ \mathsf{stack}, \_, \mathsf{cmpl}', \mathsf{c\text{-}by}', \_)}$$

$$\frac{\mathsf{return} \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{cmpl}' = \mathsf{cmpl}[i \mapsto \top]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \_, \mathsf{cmpl}, \_, \_) \xrightarrow{(aid,\, i,\, \mathsf{return})} (\mathsf{g}, \mathsf{stack}, \_, \mathsf{cmpl}', \_, \_)}$$

$$\frac{\mathsf{await}\ \mathtt{r} \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{cmpl}(\ell(r)) = \top \quad \ell' = \ell[\mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \_, \mathsf{cmpl}, \_, \_) \xrightarrow{(aid,\, i,\, \mathsf{await}(\ell(r)))} (\mathsf{g}, (i, m, \ell') \circ \mathsf{stack}, \_, \mathsf{cmpl}, \_, \_)}$$

$$\frac{\begin{array}{c} \mathsf{await}\ \mathtt{r} \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{cmpl}(\ell(r)) = \bot \quad \mathsf{w\text{-}for}' = \mathsf{w\text{-}for}[i \mapsto \ell(r)] \\ \ell' = \ell[\mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))] \end{array}}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \mathsf{pend}, \mathsf{cmpl}, \_, \mathsf{w\text{-}for}) \xrightarrow{(aid,\, i,\, \mathsf{await}(\ell(r)))} (\mathsf{g}, \mathsf{stack}, \{(i, m, \ell')\} \uplus \mathsf{pend}, \mathsf{cmpl}, \_, \mathsf{w\text{-}for}')}$$

$$\frac{\mathsf{await} * \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \ell' = \ell[\mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \_, \_, \_, \_) \xrightarrow{(aid,\, i,\, \mathsf{await}(*))} (\mathsf{g}, (i, m, \ell') \circ \mathsf{stack}, \_, \_, \_, \_)}$$

$$\frac{\mathsf{await} * \in \mathsf{inst}(\ell(\mathsf{pc})) \quad aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{w\text{-}for}' = \mathsf{w\text{-}for}[i \mapsto *] \quad \ell' = \ell[\mathsf{pc} \mapsto \mathsf{next}(\ell(\mathsf{pc}))]}{(\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \mathsf{pend}, \_, \_, \mathsf{w\text{-}for}) \xrightarrow{(aid,\, i,\, \mathsf{await}(*))} (\mathsf{g}, \mathsf{stack}, \{(i, m, \ell')\} \uplus \mathsf{pend}, \_, \_, \mathsf{w\text{-}for}')}$$

$$\frac{aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{w\text{-}for}(i) = j \quad \mathsf{cmpl}(j) = \top}{(\mathsf{g}, \epsilon, \{(i, m, \ell)\} \uplus \mathsf{pend}, \mathsf{cmpl}, \_, \mathsf{w\text{-}for}) \xrightarrow{(aid,\, i,\, \mathsf{cont})} (\mathsf{g}, (i, m, \ell), \mathsf{pend}, \mathsf{cmpl}, \_, \mathsf{w\text{-}for})}$$

$$\frac{aid \in \mathbb{A}\mathsf{id}\ \text{fresh} \quad \mathsf{w\text{-}for}(i) = *}{(\mathsf{g}, \mathsf{stack}, \{(i, m, \ell)\} \uplus \mathsf{pend}, \_, \_, \mathsf{w\text{-}for}) \xrightarrow{(aid,\, i,\, \mathsf{cont})} (\mathsf{g}, (i, m, \ell) \circ \mathsf{stack}, \mathsf{pend}, \_, \_, \mathsf{w\text{-}for})}$$

Fig. 6: Program semantics. For a function $f$, we use $f[a \mapsto b]$ to denote a function $g$ such that $g(c) = f(c)$ for all $c \neq a$ and $g(a) = b$. The function $\mathsf{inst}$ returns the instruction at some given control location while $\mathsf{next}$ gives the next instruction to execute. We use $\circ$ to denote sequence concatenation and $\mathsf{init}$ to denote the initial state of a method call.

By the definition of $\rightarrow$, every action $a \in \mathbb{A}\mathsf{ct} \setminus \{(\_, \_, \mathsf{cont})\}$ corresponds to executing some statement in the program, which is denoted by $\mathsf{S}(a)$.

An execution of $P$ is a sequence $\rho = \mathsf{ps}_0 \xrightarrow{a_1} \mathsf{ps}_1 \xrightarrow{a_2} \ldots$ of transitions starting in the initial configuration $\mathsf{ps}_0$ and leading to a configuration $\mathsf{ps}$ where the call stack and the set of pending tasks are empty. $\mathbb{C}[P]$ denotes the set of all program variable valuations included in configurations that are reached in executions of $P$. Since we are only interested in reasoning about the sequence of actions $a_1 \cdot a_2 \cdot \ldots$ labeling the transitions of an execution, we will call the latter an execution as well. The set of executions of a program $P$ is denoted by $\mathbb{E}\mathsf{x}(P)$.

**Traces.** The *trace* of execution $\rho \in \mathbb{E}\mathsf{x}(P)$ is a tuple $\mathsf{tr}(\rho) = (\rho, \mathsf{MO}, \mathsf{CO}, \mathsf{SO}, \mathsf{HB})$ of strict partial orders between the actions in $\rho$ defined in Table 1. The *method invocation order* $\mathsf{MO}$ records the order between actions in the same invocation, and the *call order* $\mathsf{CO}$ is an extension of $\mathsf{MO}$ that additionally orders actions before an invocation with respect to those inside that invocation. The *synchronous*

*happens-before order* $\mathsf{SO}$ orders the actions in an execution as if all the invocations were synchronous (even if the execution may contain asynchronous ones). It is an extension of $\mathsf{CO}$ where additionally, every action inside a callee is ordered before the actions following its invocation in the caller. The (asynchronous) *happens-before order* $\mathsf{HB}$ contains typical control-flow constraints: it is an extension of $\mathsf{CO}$ where every action $a$ inside an asynchronous invocation is ordered before the corresponding `await` in the caller, and before the actions following its invocation in the caller if $a$ precedes the first[5] `await` in $\mathsf{MO}$ (an invocation can be interrupted only when executing an `await`) or if the callee does not contain an `await` (it is synchronous). $\mathbb{T}\mathsf{r}(P)$ is the set of traces of $P$.

Table 1: Strict partial orders included in a trace. $\mathsf{CO}$, $\mathsf{SO}$, and $\mathsf{HB}$ are the smallest satisfying relations.

| | |
|---|---|
| $a_1 <_\rho a_2$ | $a_1$ occurs before $a_2$ in $\rho$ and $a_1 \neq a_2$ |
| $a_1 \sim a_2$ | $a_1 = (\_, i, \_)$ and $a_2 = (\_, i, \_)$ |
| $(a_1, a_2) \in \mathsf{MO}$ | $a_1 \sim a_2 \wedge a_1 <_\rho a_2$ |
| $(a_1, a_2) \in \mathsf{CO}$ | $(a_1, a_2) \in \mathsf{MO} \vee (a_1 = (\_, i, \mathsf{call}(j)) \wedge a_2 = (\_, j, \_))$ $\vee\ (\exists\ a_3.\ (a_1, a_3) \in \mathsf{CO} \wedge (a_3, a_2) \in \mathsf{CO})$ |
| $(a_1, a_2) \in \mathsf{SO}$ | $(a_1, a_2) \in \mathsf{CO} \vee (\exists\ a_3.\ (a_1, a_3) \in \mathsf{SO} \wedge (a_3, a_2) \in \mathsf{SO})$ $\vee\ (a_1 = (\_, j, \_) \wedge a_2 = (\_, i, \_) \wedge \exists\ a_3 = (\_, i, \mathsf{call}(j)).\ a_3 <_\rho a_2)$ |
| $(a_1, a_2) \in \mathsf{HB}$ | $(a_1, a_2) \in \mathsf{CO} \vee (\exists\ a_3.\ (a_1, a_3) \in \mathsf{HB} \wedge (a_3, a_2) \in \mathsf{HB})$ $\vee\ (\ a_1 = (\_, j, \_) \wedge a_2 = (\_, i, \_) \wedge \exists\ a_3 = (\_, i, \mathsf{await}(j)).\ a_3 <_\rho a_2\ )$ $\vee\ (\ a_1 = (\_, j, \mathsf{await}(i'))$ is the first await in $j \wedge$ $a_2 = (\_, i, \_) \wedge \exists\ a_3 = (\_, i, \mathsf{call}(j)).\ a_3 <_\rho a_2\ )$ $\vee\ (\ a_1 = (\_, j, \_) \wedge \nexists\ (\_, j, \mathsf{await}(\_)) \in \rho \wedge$ $a_2 = (\_, i, \_) \wedge \exists\ a_3 = (\_, i, \mathsf{call}(j)).\ a_3 <_\rho a_2\ )$ |

On the right of Fig. 1, we show a trace where two statements (represented by the corresponding lines numbers) are linked by a dotted arrow if the corresponding actions are related by $\mathsf{MO}$, a dashed arrow if the corresponding actions are related by $\mathsf{CO}$ but not by $\mathsf{MO}$, and a solid arrow if the corresponding actions are related by the $\mathsf{HB}$ but not by $\mathsf{CO}$.

## 3 Synthesizing Asynchronous Programs

Given a synchronous program $P$ and a subset of *base* methods $L \subseteq P$, our goal is to synthesize *all* asynchronous programs $P_a$ that are equivalent to $P$ and that are obtained by substituting every method in $L$ with an equivalent *asynchronous* version. The base methods are considered to be models of standard library calls (e.g., IO operations) and asynchronous versions are defined by inserting `await` $*$ statements in their body. We use $P[L]$ to emphasize a subset of base methods $L$ in a program $P$. Also, we call $L$ a *library*. A library is called (a)synchronous when all methods are (a)synchronous.

[5] Code in between two awaits can execute before or after the control is returned to the caller, depending on whether the first awaited task finished or not.

**Asynchronizations of a synchronous program.** Let $P[L]$ be a synchronous program, and $L_a$ a set of asynchronous methods obtained from those in $L$ by inserting at least one `await` $*$ statement in their body (and adding the keyword `async`). Each method in $L_a$ corresponds to a method in $L$ with the same name, and vice-versa. $P_a[L_a]$ is called an *asynchronization* of $P[L]$ with respect to $L_a$ if it is a syntactically correct program obtained by replacing the methods in $L$ with those in $L_a$ and adding `await` statements as necessary. More precisely, let $L^* \subseteq P$ be the set of all methods of $P$ that transitively call methods of $L$. Formally, $L^*$ is the smallest set of methods that includes $L$ and satisfies the following: if a method $m$ calls $m' \in L^*$, then $m \in L^*$. Then, $P_a[L_a]$ is an *asynchronization* of $P[L]$ w.r.t. $L_a$ if it is obtained from $P$ as follows:

```
method m {          async method m {      async method m {
  r1 = call m1;       r1 = call m1;         r1 = call m1;
                      await r1;
  r2 = x;             r2 = x;               r2 = x;
            }                     }               await r1;       }
method m1 {         async method m1 {     async method m1 {
                      await *               await *
  retVal = x;         retVal = x;           retVal = x;
  x = input;          x = input;            x = input;
  return;  }          return;       }       return;       }
```

Fig. 7: A program and its asynchronizations.

- Each method in $L$ is replaced with the corresponding method from $L_a$.
- All methods in $L^* \setminus L$ are declared as asynchronous (because every call to an asynchronous method is followed by an `await` and any method using `await` must be asynchronous).
- For each invocation $r :=$ call $m$ of $m \in L^*$, add `await` statements await $r$ satisfying the well-formedness syntactic constraints described in Section 2.

Fig. 7 lists a synchronous program and its two asynchronizations, where $L = \{m1\}$ and $L^* = \{m, m1\}$. Asynchronizations differ only in the await placement.

$\mathsf{Asy}[P, L, L_a]$ is the set of all asynchronizations of $P[L]$ w.r.t. $L_a$. The *strong* asynchronization $\mathsf{strongAsy}[P, L, L_a]$ is an asynchronization where every `await` *immediately* follows the matching call. It reaches exactly the same set of program variable valuations as $P$.

**Problem definition.** We investigate the problem of enumerating *all* asynchronizations of a given program w.r.t. a given asynchronous library, which are *sound*, in the sense that they do not admit data races. Two actions $a_1$ and $a_2$ in a trace $\tau = (\rho, \mathsf{MO}, \mathsf{CO}, \mathsf{SO}, \mathsf{HB})$ are *concurrent* if $(a_1, a_2) \notin \mathsf{HB}$ and $(a_2, a_1) \notin \mathsf{HB}$.

An ansynchronous program $P_a$ *admits a data race* $(a_1, a_2)$, where $(a_1, a_2) \in \mathsf{SO}$, if $a_1$ and $a_2$ are two concurrent actions of a trace $\tau \in \mathbb{T}\mathsf{r}(P_a)$, and $a_1$ and $a_2$ are read or write accesses to the same program variable $x$, and at least one of them is a write. We write data races as ordered pairs w.r.t. $\mathsf{SO}$ to simplify the definition of the algorithms in the next sections. Also, note that traces of *synchronous* programs can *not* contain concurrent actions, and therefore they do not admit data races. $\mathsf{strongAsy}[P, L, L_a]$ does not admit data races as well.

$P_a[L_a]$ is called *sound* when it does not admit data races. The absence of data races implies equivalence to the original program, in the sense of reaching the same set of configurations (program variable valuations).

**Definition 1.** *For a synchronous program $P[L]$ and asynchronous library $L_a$, the* asychronization synthesis problem *asks to enumerate all sound asynchronizations in* $\mathsf{Asy}[P, L, L_a]$.

## 4 Enumerating Sound Asynchronizations

We present an algorithm for solving asynchronization synthesis, which relies on a partial order between asynchronizations that guides the enumeration of possible solutions. The partial order takes into account the distance between calls and corresponding awaits. Fig. 8 pictures the partial order for asynchronizations of the program on the left of Fig. 1. Each asynchronization is written as a vector of distances, the first (second) element is the number of statements between `await t1` (`await t`) and the matching call (we count only statements that appear in the sequential program). The edges connect comparable elements, smaller elements being below bigger elements. The asynchronization on the middle of Fig. 1 corresponds to the vector $(1, 1)$. The highlighted elements constitute the set of all sound asynchronizations. The strong asynchronization corresponds to the vector $(0, 0)$.

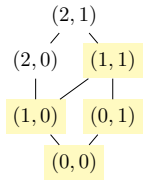

Fig. 8

Formally, an `await` statement $s_w$ in a method $m$ of an asynchronization $P_a[L_a] \in \mathsf{Asy}[P, L, L_a]$ *covers* a read/write statement $s$ in $P$ if there exists a path in the CFG of $m$ from the call statement matching $s_w$ to $s_w$ that contains $s$. The set of statements covered by an await $s_w$ is denoted by $\mathsf{Cover}(s_w)$. We compare asynchronizations in terms of sets of statements covered by awaits that match the same call from the synchronous program $P[L]$. Since asynchronizations are obtained by adding `await`s, every call in asynchronization $P_a[L_a] \in \mathsf{Asy}[P, L, L_a]$ corresponds to a *fixed* call in $P[L]$. Therefore, for two asynchronizations $P_a, P'_a \in \mathsf{Asy}[P, L, L_a]$, $P_a$ is *smaller* than $P'_a$, denoted by $P_a \leq P'_a$, iff for every `await` $s_w$ in $P_a$, there exists an `await` $s'_w$ in $P'_a$ that matches the same call as $s_w$, such that $\mathsf{Cover}(s_w) \subseteq \mathsf{Cover}(s'_w)$. For example, the two asynchronous programs in Fig. 7 are ordered by $\leq$ since $\mathsf{Cover}(\mathsf{await}\ r1) = \{\}$ in the first and $\mathsf{Cover}(\mathsf{await}\ r1) = \{\texttt{r2 = x}\}$ in the second. Note that the strong asynchronization is smaller than every other asynchronization. Also, note that $\leq$ has a unique maximal element that is called the weakest asynchronization and denoted by $\mathsf{wkAsy}[P, L, L_a]$. In Fig. 8, the weakest asynchronization corresponds to the vector $(2, 1)$.

In the following, we say *moving an await down (resp., up)* when moving the await further away from (resp. closer to) the matching call while preserving well-formedness conditions in Section 2. Further away or closer to means increasing or decreasing the set of statements that are covered by the await. For instance, if an await $s_w$ in a program $P_a$ is preceded by a while loop, then *moving it up* means moving it before the whole loop and not inside the loop body. Otherwise, the third well-formedness condition would be violated.

**Relative Maximality.** A crucial property of this partial order is that for every asynchronization $P_a$, there exists a *unique* maximal asynchronization that is smaller than $P_a$ and that is sound. Formally, an asynchronization $P'_a$ is called a

**Algorithm 1** An algorithm for enumerating all sound asynchronizations (these asynchronizations are obtained as a result of the **output** instruction). MaxRel returns the maximal asynchronization of $P$ relative to $P_a$

```
1: procedure AsySyn(P_a, s_w)
2:     P'_a ← MaxRel(P_a);
3:     output P'_a;
4:     P ← ImPred(P'_a, s_w);
5:     for each (P''_a, s''_w) ∈ P
6:         AsySyn(P''_a, s''_w);
```

*maximal asynchronization of $P$ relative to $P_a$* if (1) $P'_a \leq P_a$, $P'_a$ is sound, and (2) $\forall\ P''_a \in \mathsf{Asy}[P, L, L_a]$. $P''_a$ is sound and $P''_a \leq P_a \Rightarrow P''_a \leq P'_a$.

**Lemma 1.** *Given an asynchronization $P_a \in \mathsf{Asy}[P, L, L_a]$, there exists a unique program $P'_a$ that is a maximal asynchronization of $P$ relative to $P_a$.*

The asynchronization $P'_a$ exists because the bottom element of $\leq$ is sound. To prove uniqueness, assume by contradiction that there exist two incomparable maximal asynchronizations $P^1_a$ and $P^2_a$ and select the first await $s^1_w$ w.r.t. the control-flow of the sequential program that is placed in different positions in the two programs. Assume that $s^1_w$ is closer to its matching call in $P^1_a$. Then, we move $s^1_w$ in $P^1_a$ further away from its matching call to the same position as in $P^2_a$. This modification does not introduce data races since $P^2_a$ is data race free. Thus, the resulting program is data race free, bigger than $P^1_a$, and smaller than $P_a$ w.r.t. $\leq$ contradicting the fact that $P^1_a$ is a maximal asynchronization.

### 4.1 Enumeration Algorithm

Our algorithm for enumerating all sound asynchronizations is given in Algorithm 1 as a recursive procedure AsySyn that we describe in two phases.

First, ignore the second argument of AsySyn (in blue), which represents an `await` statement. For an asynchronization $P_a$, AsySyn outputs *all* sound asynchronizations that are smaller than $P_a$. It uses MaxRel to compute the maximal asynchronization $P'_a$ of $P$ relative to $P_a$, and then, calls itself recursively for all immediate predecessors of $P'_a$. AsySyn outputs all sound asynchronizations of $P$ when given as input the weakest asynchronization of $P$.

Recursive calls on immediate predecessors are necessary because the set of sound asynchronizations is not downward-closed w.r.t. $\leq$. For instance, the asynchronization on the right of Fig. 9 is an immediate predecessor of the sound asynchronization on the left but it has a data race on $x$.

```
async method m {          async method m {
  r1 = call m1;             r1 = call m1;
  r2 = x;                   r2 = x;
  await r1;     }           await r1;     }
async method m1 {         async method m1 {
  r3 = call m2;             r3 = call m2;
  x = x + 1;                await r3;
  await r3;     }           x = x + 1;    }
async method m2 {         async method m2 {
  await *                   await *
  retVal = input;           retVal = input;
  return;       }           return;       }
```

Fig. 9: Asynchronizations.

The delay complexity of this algorithm remains exponential in general, since a sound asynchronization may be outputted multiple times. Asynchronizations are only partially ordered by $\leq$ and

different chains of recursive calls starting in different immediate predecessors may end up outputting the same solution. For instance, for the asynchronizations in Fig. 8, the asynchronization $(0,0)$ will be outputted twice because it is an immediate predecessor of both $(1,0)$ and $(0,1)$.

To avoid this redundancy, we use a refinement of the above that *restricts* the set of immediate predecessors available for a (recursive) call of AsySyn. This is based on a *strict total order* $\prec_w$ between `await`s in a program $P_a$ that follows a topological ordering of its inter-procedural CFG, i.e., if $s_w$ occurs before $s'_w$ in the body of a method $m$, then $s_w \prec_w s'_w$, and if $s_w$ occurs in a method $m$ and $s'_w$ occurs in a method $m'$ s.t. $m$ (indirectly) calls $m'$, then $s_w \prec_w s'_w$. Therefore, AsySyn takes an await statement $s_w$ as a second parameter, which is initially the maximal element w.r.t. $\prec_w$, and it calls itself only on immediate predecessors of a solution obtained by *moving up* an await $s''_w$ *smaller than or equal to* $s_w$ w.r.t. $\prec_w$. The recursive call on that predecessor will receive as input $s''_w$. Formally, this relies on a function $\mathsf{ImPred}$ that returns pairs of immediate predecessors and await statements defined as follows:

$$\mathsf{ImPred}(P'_a, s_w) = \{(P''_a, s''_w) : P''_a < P'_a \text{ and } \forall P'''_a \in \mathsf{Asy}[P,L,L_a].P'''_a < P'_a \Rightarrow P'''_a \leq P''_a \text{ and } s''_w \preceq_w s_w \text{ and } P''_a \in P'_a \uparrow s''_w \}$$

($P'_a \uparrow s''_w$ is the set of asynchronizations obtained from $P'_a$ by changing *only* the position of $s''_w$, moving it up w.r.t. the position in $P'_a$). For instance, looking at immediate predecessors of $(1,1)$ in Fig. 8, $(0,1)$ is obtained by moving the *first* await in $\prec_w$. Therefore, the recursive call on $(0,1)$ computes the maximal asynchronization relative to $(0,1)$, which is $(0,1)$, and stops ($\mathsf{ImPred}$ returns $\emptyset$ because the input $s_w$ is the minimal element of $\prec_w$, and already immediately after the call). Its immediate predecessor is explored when recursing on $(1,0)$.

Algorithm 1 outputs all sound asynchronizations because after having computed a maximal asynchronization $P'_a$ in a recursive call with parameter $s_w$, any smaller sound asynchronization is smaller than a predecessor in $\mathsf{ImPred}(P'_a, s_w)$. Also, it can not output the same asynchonization twice. Let $P^1_a$ and $P^2_a$ be two predecessors in $\mathsf{ImPred}(P'_a, s_w)$ obtained by moving up the awaits $s^1_w$ and $s^2_w$, respectively, and assume that $s^1_w \prec_w s^2_w$. Then, all solutions computed in the recursive call on $P^1_a$ will have $s^2_w$ placed as in $P'_a$ while all the solutions computed in the recursive call on $P^2_a$ will have $s^2_w$ closer to the matching call. Therefore, the sets of solutions computed in these two recursion branches are distinct.

**Theorem 1.** AsySyn$(\mathsf{wkAsy}[P,L,L_a], s_w)$ *where* $s_w$ *is the maximal await in* $\mathsf{wkAsy}[P,L,L_a]$ *w.r.t.* $\prec_w$ *outputs all sound asynchronizations of* $P[L]$ *w.r.t.* $L_a$.

The delay complexity of Algorithm 1 is polynomial time modulo an oracle that returns a maximal asynchronization relative to a given one. In the next section, we show that the latter problem can be reduced in polynomial time to the reachability problem in sequential programs.

## 5 Computing Maximal Asynchronizations

In this section, we present an implementation of the procedure MaxRel that relies on a reachability oracle. In particular, we first describe an approach for computing the maximal asynchronization relative to a given asynchronization $P_a$, which can be seen as a way of repairing $P_a$ so that it becomes data-race free. Intuitively, we repeatedly eliminate data races in $P_a$ by moving certain `await` statements closer to the matching calls. The data races in $P_a$ (if any) are enumerated in a certain order that prioritizes data races between actions that occur first in executions of the original synchronous program. This order allows to avoid superfluous repair steps.

### 5.1 Data Race Ordering

An action $a$ representing a read/write access in a trace $\tau$ of an asynchronization $P_a$ of $P$ is *synchronously reachable* if there is an action $a'$ in a trace $\tau'$ of $P$ that represents the same statement, i.e., $\mathsf{S}(a) = \mathsf{S}(a')$. It can be proved that any trace of an asynchronization contains a data race if it contains a data race between two synchronously reachable actions (see Appendix C). In the following, we focus on data races between actions that are synchronously reachable.

We define an order between such data races based on the order between actions in executions of the original synchronous program $P$. This order relates data races in possibly different executions or asynchronizations of $P$, which is possible because each action in a data race corresponds to a statement in $P$.

For two read/write statements $s$ and $s'$, $s \prec s'$ denotes the fact that there is an execution of $P$ in which the *first* time $s$ is executed occurs before the *first* time $s'$ is executed. For two actions $a$ and $a'$ in an execution/trace of an asynchronization, generated by two read/write statements $s = \mathsf{S}(a)$ and $s' = \mathsf{S}(a')$, $a \prec_{\mathsf{SO}} a'$ holds if $s \prec s'$ and either $s' \not\prec s$ or $s'$ is reachable from $s$ in the interprocedural[6] control-flow graph of $P$ without taking any back edge[7]. For a *deterministic* synchronous program (admitting a single execution), $a \prec_{\mathsf{SO}} a'$ iff $\mathsf{S}(a) \prec \mathsf{S}(a')$. For non-deterministic programs, when $\mathsf{S}(a)$ and $\mathsf{S}(a')$ are contained in a loop body, it is possible that $\mathsf{S}(a) \prec \mathsf{S}(a')$ and $\mathsf{S}(a') \prec \mathsf{S}(a)$. In this case, we use the control-flow order to break the tie between $a$ and $a'$.

The order between data races corresponds to the colexicographic order induced by $\prec_{\mathsf{SO}}$. This is a partial order since actions may originate from different control-flow paths and are incomparable w.r.t. $\prec_{\mathsf{SO}}$.

**Definition 2 (Data Race Order).** *Given two races $(a_1, a_2)$ and $(a_3, a_4)$ admitted by (possibly different) asynchronizations of a synchronous program $P$, we have that $(a_1, a_2) \prec_{\mathsf{SO}} (a_3, a_4)$ iff $a_2 \prec_{\mathsf{SO}} a_4$, or $a_2 = a_4$ and $a_1 \prec_{\mathsf{SO}} a_3$.*

[6] The interprocedural graph is the union of the control-flow graphs of each method along with edges from call sites to entry nodes, and from exit nodes to return sites.

[7] A back edge points to a block that has already been met during a depth-first traversal of the control-flow graph, and corresponds to loops.

Repairing a minimal data race $(a_1, a_2)$ w.r.t. $\prec_{\mathsf{SO}}$ removes any other data race $(a_1, a_4)$ with $(a_2, a_4) \in \mathsf{HB}$ (note that we cannot have $(a_4, a_2) \notin \mathsf{HB}$ since $a_2 \prec_{\mathsf{SO}} a_4$). The repair will enforce that $(a_1, a_2) \in \mathsf{HB}$ which implies that $(a_1, a_4) \in \mathsf{HB}$.

### 5.2 Repairing Data Races

Repairing a data race $(a_1, a_2)$ reduces to modifying the position of a certain `await`. We consider only repairs where `await`s are moved up (closer to the matching call). The "completeness" of this set of repairs follows from the particular order in which we enumerate data races.

Let $s_1$ and $s_2$ be the statements generating $a_1$ and $a_2$. In general, there exists a method $m$ that (transitively) calls another asynchronous method $m1$ that contains $s_1$ and before awaiting for $m1$ it (transitively) calls a method $m2$ that executes $s_2$. This is pictured in Fig. 10. It is also possible that $m$ itself contains $s_2$ (see the program on the right of Fig. 7). The repair consists in moving the await for $m1$ before the call to $m2$ since this implies that $s_1$ will always execute before $s_2$ (and the corresponding actions are related by happens-before).

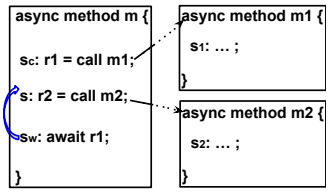


Fig. 10: A data race repair.

Formally, any two racing actions have a common ancestor in the call order $\mathsf{CO}$ which is a call action. The least common ancestor of $a_1$ and $a_2$ in $\mathsf{CO}$ among call actions is denoted by $\mathsf{LCA}_{\mathsf{CO}}(a_1, a_2)$. In Fig. 10, it corresponds to the call statement $s_c$. More precisely, $\mathsf{LCA}_{\mathsf{CO}}(a_1, a_2)$ is a call action $a_c = (\_, i, \mathsf{call}(j))$ s.t. $(a_c, a_1) \in \mathsf{CO}$, $(a_c, a_2) \in \mathsf{CO}$, and for each other call action $a'_c$, if $(a_c, a'_c) \in \mathsf{CO}$ then $(a'_c, a_1) \notin \mathsf{CO}$. This call action represents an asynchronous call for which the matching await $s_w$ must move to repair the data race. The await should be moved before the last statement in the same method generating an action which precedes $a_2$ in the reflexive closure of call order (statement $s$ in Fig. 10). This way every statement that follows $s_c$ in call order will be executed before $s$ and before any statement which succeeds $s$ in call order, including $s_2$. Note that moving the await $s_w$ anywhere after $s$ will not affect the concurrency between $a_1$ and $a_2$.

The pair $(s_c, s)$ is called the *root cause* of the data race $(a_1, a_2)$. We denote by $\mathsf{RDR}(P_a, s_c, s)$ the maximal asynchronization $P'_a$ smaller than $P_a$ w.r.t. $\leq$, s.t. no await statement matching $s_c$ occurs after $s$ on a CFG path.

### 5.3 A Procedure for Computing Maximal Asynchronizations

Given an asynchronization $P_a$, the procedure MaxRel in Algorithm 2 computes the maximal asynchronization relative to $P_a$ by repairing data races iteratively until the program becomes data race free. The sub-procedure RCMinDR$(P'_a)$ computes the root cause of a minimal data race $(a_1, a_2)$ of $P'_a$ w.r.t. $\prec_{\mathsf{SO}}$ such that the two actions are synchronously reachable. If $P'_a$ is data race free, then RCMinDR$(P'_a)$ returns $\perp$. The following theorem states the correctness of MaxRel.

**Algorithm 2** The procedure MaxRel to find the maximal asynchronization of $P$ relative to $P_a$.

```
1: procedure MaxRel(P_a)
2:     P'_a ← P_a
3:     root ← RCMinDR(P'_a)
4:     while root ≠ ⊥
5:         P'_a ← RDR(P'_a, root)
6:         root ← RCMinDR(P'_a)
7:     return P'_a
```

**Theorem 2.** *Given an asynchronization $P_a \in \mathsf{Asy}[P, L, L_a]$, MaxRel$(P_a)$ returns the maximal asynchronization of $P$ relative to $P_a$.*

MaxRel$(P_a)$ repairs a number of data races which is linear in the size of the input. Indeed, each repair results in moving an await closer to the matching call and before at least one more statement from the original program $P$.

The problem of computing root causes of minimal data races is reducible to reachability (assertion checking) in sequential programs. This reduction builds on a program instrumentation for checking if there exists a data race that involves two given statements $(s_1, s_2)$ that are reachable in an executions of $P$.This instrumentation is used in an iterative process where pairs of statements are enumerated according to the colexicographic order induced by $\prec$. For lack of space, we present only the main ideas of the instrumentation (see Appendix D). The instrumentation simulates executions of an asynchronization $P_a$ using nondeterministic synchronous code where methods may be only partially executed (modeling `await` interruptions). Immediately after executing $s_1$, the current invocation $t_1$ is interrupted (by executing a `return` added by the instrumentation). The active invocations that transitively called $t_1$ are also interrupted when reaching an `await` for an invocation in this call chain (the other invocations are executed until completion as in the synchronous semantics). When reaching $s_2$, if $s_1$ has already been executed and at least one invocation has been interrupted, which means that $s_1$ is concurrent with $s_2$, then the instrumentation stops with an assertion violation. The instrumentation also computes the root cause of the data race using additional variables for tracking call dependencies.

## 6 Asymptotic Complexity of Asynchronization Synthesis

We state the complexity of the asynchronization synthesis problem. Algorithm 1 shows that the delay complexity of this problem is polynomial-time in the number of statements in input program modulo the complexity of computing a maximal asynchronization, which Algorithm 2 shows to be polynomial-time reducible to reachability in sequential programs. Since the reachability problem is PSPACE-complete for finite-state sequential programs [15], we get the following:

**Theorem 3.** *The output complexity[8] and delay complexity of the asynchronization synthesis problem is polynomial time modulo an oracle for reachability in sequential programs, and PSPACE for finite-state programs.*

[8] Note that all asynchronizations can be enumerated with polynomial space.

This result is optimal, i.e., checking whether there exists a sound asynchronization which is different from the trivial strong synchronization is PSPACE-hard (follows from a reduction from the reachability problem). See Appendices D and E for the detailed formal proofs.

## 7 Asynchronization Synthesis Using Data-Flow Analysis

In this section, we present a refinement of Algorithm 2 that relies on a bottom-up inter-procedural data flow analysis. The analysis is used to compute maximal asynchronizations for abstractions of programs where every Boolean condition (in if-then-else or while statements) is replaced with the non-deterministic choice $*$, and used as an implementation of MaxRel in Algorithm 1.

For a program $P$, we define an abstraction $P^{\#}$ where every conditional `if` $\langle le \rangle$ $\{S_1\}$ `else` $\{S_2\}$ is rewritten to `if` $*$ $\{S_1\}$ `else` $\{S_2\}$, and every `while` $\langle le \rangle$ $\{S\}$ is rewritten to `if` $*$ $\{S\}$. Besides adding the non-deterministic choice $*$, loops are unrolled exactly once. Every asynchronization $P_a$ of $P$ corresponds to an abstraction $P_a^{\#}$ obtained by applying exactly the same rewriting. $P^{\#}$ is a sound abstraction of $P$ in terms of sound asynchronizations it admits. Unrolling loops once is sound because every asynchronous call in a loop iteration should be awaited for in the same iteration (see the syntactic constraints in Section 2).

**Theorem 4.** *If $P_a^{\#}$ is a sound asynchronization of $P^{\#}$ w.r.t. $L_a$, then $P_a$ is a sound asynchronization of $P$ w.r.t. $L_a$.*

The procedure for computing maximal asynchronizations of $P^{\#}$ relative to a given asynchronization $P_a^{\#}$ traverses methods of $P_a^{\#}$ in a bottom-up fashion, detects data races using summaries of read/write accesses computed using a straightforward data-flow analysis, and repairs data races using the schema presented in Section 5.2. Applying this procedure to a real programming language requires an alias analysis to detect statements that may access the same memory location (this is trivial in our language which is used to simplify the exposition).

We consider an enumeration of methods called *bottom-up order*, which is the reverse of a topological ordering of the call graph[9]. For each method $m$, let $\mathcal{R}(m)$ be the set of program variables that $m$ can read, which is defined as the union of $\mathcal{R}(m')$ for every method $m'$ called by $m$ and the set of program variables read in statements in the body of $m$. The set of variables $\mathcal{W}(m)$ that $m$ can write is defined in a similar manner. We define $\mathsf{RW\text{-}var}(m) = (\mathcal{R}(m), \mathcal{W}(m))$. We extend the notation $\mathsf{RW\text{-}var}$ to statements as follows: $\mathsf{RW\text{-}var}(\langle r \rangle := \langle x \rangle) = (\{x\}, \emptyset)$, $\mathsf{RW\text{-}var}(\langle x \rangle := \langle le \rangle) = (\emptyset, \{x\})$, $\mathsf{RW\text{-}var}(r := \texttt{call}\ m) = \mathsf{RW\text{-}var}(m)$, and $\mathsf{RW\text{-}var}(s) = (\emptyset, \emptyset)$, for any other type of statement $s$. Also, let $\mathsf{CRW\text{-}var}(m)$ be the set of read or write accesses that $m$ can do and that can be concurrent with accesses that a caller of $m$ can do after calling $m$. These correspond to read/write statements that follow an `await` in $m$, or to accesses in $\mathsf{CRW\text{-}var}(m')$ for a method $m'$ called by $m$. These sets of accesses can be computed using the

[9] The nodes of the call graph are methods and there is an edge from a method $m_1$ to a method $m_2$ if $m_1$ contains a call statement that calls $m_2$.

following data-flow analysis: for all methods $m \in P_a^{\#}$ in bottom-up order, and for each statement $s$ in the body of $m$ from begin to end,

- if $s$ is a call to $m'$ and $s$ is *not* reachable from an `await` in the CFG of $m$
  - $\mathsf{CRW\text{-}var}(m) \leftarrow \mathsf{CRW\text{-}var}(m) \cup \mathsf{CRW\text{-}var}(m')$
- if $s$ is reachable from an `await` statement in the CFG of $m$
  - $\mathsf{CRW\text{-}var}(m) \leftarrow \mathsf{CRW\text{-}var}(m) \cup \mathsf{RW\text{-}var}(s)$

We use $(\mathcal{R}_1, \mathcal{W}_1) \bowtie (\mathcal{R}_2, \mathcal{W}_2)$ to denote the fact that $\mathcal{W}_1 \cap (\mathcal{R}_2 \cup \mathcal{W}_2) \neq \emptyset$ or $\mathcal{W}_2 \cap (\mathcal{R}_1 \cup \mathcal{W}_1) \neq \emptyset$ (i.e., a conflict between read/write accesses). We define the procedure MaxRel$^{\#}$ that given an asynchronization $P_a^{\#}$ works as follows:

- for all methods $m \in P_a^{\#}$ in bottom-up order, and for each statement $s$ in the body of $m$ from begin to end,
  - if $s$ occurs between $r :=$ `call` $m'$ and `await` $r$ (for some $m'$), and $\mathsf{RW\text{-}var}(s) \bowtie \mathsf{CRW\text{-}var}(m')$, then $P_a^{\#} \leftarrow \mathsf{RDR}(P_a^{\#}, r :=$ `call` $m', s)$
- return $P_a^{\#}$

**Theorem 5.** *The procedure* MaxRel$^{\#}(P_a^{\#})$ *returns a maximal asynchronization relative to* $P_a^{\#}$.

Since MaxRel$^{\#}$ is based on a single bottom-up traversal of the call graph of the input asynchronization $P_a^{\#}$ we get the following result.

**Theorem 6.** *The delay complexity of the asynchronization synthesis problem restricted to abstracted programs* $P^{\#}$ *is polynomial time.*

## 8 Experimental Evaluation

We present an empirical evaluation of our asynchronization synthesis approach, where maximal asynchronizations are computed using the data-flow analysis in Section 7. Our benchmark consists mostly of asynchronous C# programs from open-source GitHub projects. We evaluate the effectiveness in reproducing the original program as an asynchronization of a program where asynchronous calls are reverted to synchronous calls, along with other sound asynchronizations.

**Implementation.** We developed a prototype tool that uses the Roslyn .NET compiler platform [26] to construct CFGs for methods in a C# program. This prototype supports C# programs written in static single assignment (SSA) form that include basic conditional/looping constructs and async/await as concurrency primitives. Note that object fields are interpreted as program variables in the terminology of §2 (data races concern accesses to object fields). It assumes that alias information is provided apriori; these constraints can be removed in the future with more engineering effort. In general, our synthesis procedure is compatible with any sound alias analysis. The precision of this analysis impacts only the set (number) of asynchronizations outputted by the procedure (a more precise analysis may lead to more sound asynchronizations).

The tool takes as input a possibly asynchronous program, and a mapping between synchronous and asynchronous variations of base methods in this program. It reverts every asynchronous call to a synchronous call, and it enumerates sound asynchronizations of the obtained program (using Algorithm 1).

Table 2: Empirical results. Syntactic characteristics of input programs: lines of code (loc), number of methods (m), number of method calls (c), number of asynchronous calls (ac), number of awaits that *could* be placed at least one statement away from the matching call ($\text{await}_{\#}$). Data concerning the enumeration of asynchronizations: number of awaits that *were* placed at least one statement away from the matching call (await), number of races discovered and repaired (races), number of statements that the awaits in the maximal asynchronization are covering *more than* in the input program (cover), number of computed asynchronizations (async), and running time (t).

| Program | loc | m | c | ac | $\text{await}_{\#}$ | await | races | cover | async | t(s) |
|---|---|---|---|---|---|---|---|---|---|---|
| SyntheticBenchmark-1 | 77 | 3 | 6 | 5 | 4 | 4 | 5 | 0 | 9 | 1.4 |
| SyntheticBenchmark-2 | 115 | 4 | 12 | 10 | 6 | 3 | 3 | 0 | 8 | 1.4 |
| SyntheticBenchmark-3 | 168 | 6 | 16 | 13 | 9 | 7 | 4 | 0 | 128 | 1.5 |
| SyntheticBenchmark-4 | 171 | 6 | 17 | 14 | 10 | 8 | 5 | 0 | 256 | 1.9 |
| SyntheticBenchmark-5 | 170 | 6 | 17 | 14 | 10 | 8 | 9 | 0 | 272 | 2 |
| Azure-Remote | 520 | 10 | 14 | 5 | 0 | 0 | 0 | 0 | 1 | 2.2 |
| Azure-Webjobs | 190 | 6 | 14 | 6 | 1 | 1 | 0 | 1 | 3 | 1.6 |
| FritzDectCore | 141 | 7 | 11 | 8 | 1 | 1 | 0 | 1 | 2 | 1.6 |
| MultiPlatform | 53 | 2 | 6 | 4 | 2 | 2 | 0 | 2 | 4 | 1.1 |
| NetRpc | 887 | 13 | 18 | 11 | 4 | 1 | 3 | 0 | 3 | 2 |
| TestAZureBoards | 43 | 3 | 3 | 3 | 0 | 0 | 0 | 0 | 1 | 1.5 |
| VBForums-Viewer | 275 | 7 | 10 | 7 | 3 | 2 | 1 | 1 | 6 | 1.8 |
| Voat | 178 | 3 | 5 | 5 | 2 | 1 | 1 | 1 | 3 | 1.2 |
| WordpressRESTClient | 133 | 3 | 10 | 8 | 4 | 2 | 1 | 0 | 4 | 1.7 |
| ReadFile-Stackoverflow | 47 | 2 | 3 | 3 | 1 | 0 | 1 | 0 | 1 | 1.5 |
| UI-Stackoverflow | 50 | 3 | 4 | 4 | 3 | 3 | 3 | 0 | 12 | 1.5 |

**Benchmark.** Our evaluation uses a benchmark listed in Table 2, which contains 5 synthetic examples (variations of the program in Fig. 1), 9 programs extracted from open-source C# GitHub projects (their name is a prefix of the repository name), and 2 programs inspired by questions on `stackoverflow.com` about async/await in C# (their name ends in Stackoverflow). Overall, there are 13 base methods involved in computing asynchronizations of these programs (having both synchronous and asynchronous versions), coming from 5 C# libraries (*System.IO*, *System.Net*, *Windows.Storage*, *Microsoft.WindowsAzure.Storage*, and *Microsoft.Azure.Devices*). They are modeled as described in § 2.

**Evaluation.** The last five columns of Table 2 list data concerning the application of our tool. The column async lists the number of outputted sound asynchronizations. In general, the number of asynchronizations depends on the number of invocations (column ac) and the size of the code blocks between an invocation and the instruction using its return value (column $\text{await}_{\#}$ gives the number of non-empty blocks). The number of *sound* asynchronizations depends roughly, on how many of these code blocks are racing with the method body. These asynchronizations contain `await`s that are at a non-zero distance from the matching call (non-zero values in column await) and for many Github programs, this distance is bigger than in the original program (non-zero values in column cover). This shows that we are able to increase the distances between `await`s and their matching calls for those programs. The distance between `await`s and matching calls in maximal asynchronizations of non synthetic benchmarks is 1.27 statements on average. A statement representing a method call is counted as one

independently of the method's body size. With a single level of inlining, the number of statements becomes 2.82 on average. However, these statements are again, mostly IO calls (access to network or disk) or library calls (string/bytes formatting methods) whose execution time is not negligible. The running times for the last three synthetic benchmarks show that our procedure is scalable when programs have a large number of sound asynchronizations.

With few exceptions, each program admits multiple sound asynchronizations (values in column async bigger than one), which makes the focus on the delay complexity relevant. This leaves the possibility of making a choice based on other criteria, e.g., performance metrics. As shown by the examples in Fig. 2, their performance can be derived only dynamically (by executing them). These results show that our techniques have the potential of becoming the basis of a refactoring tool allowing programmers to improve their usage of the async/await primitives. The artifacts are available in a GitHub repository [2].

## 9 Related Work

There are many works on synthesizing or repairing concurrent programs in the standard multi-threading model, e.g., automatic parallelization in compilers [1, 6, 18], or synchronization synthesis [10, 11, 23, 30, 29, 5, 9, 17]. We focus on the use of async/await which poses specific challenges not covered in these works.

Our semantics without `await` $*$ instructions is equivalent to the semantics defined in [3, 27]. But, to simplify the exposition, we consider a more restricted programming language. For the modeling of asynchronous IO operations, we follow [3] with the restriction that the code following an `await` $*$ is executed atomically. This is sound when focusing on data-race freedom because even if executed atomically, any two instructions from different asynchronous IO operations (following `await` $*$) are not happens-before related.

**Program Refactoring.** Program refactoring tools have been proposed for converting C# programs using explicit callbacks into async/await programs [24] or Android programs using AsyncTask into programs that use IntentService [21]. The C# tool [24], which is the closest to our work, makes it possible to repair misusage of async/await that might result in deadlocks. This tool cannot modify procedure calls to be asynchronous as in our work. A static analysis based technique for refactoring JavaScript programs is proposed in [16]. As opposed to our work, this refactoring technique is unsound in general. It requires that programmers review the refactoring for correctness, which is error-prone. Also, in comparison to [16], we carry a formal study of the more general problem of finding all sound asynchronizations and investigate its complexity.

**Data Race Detection.** Many works study dynamic data race detection using happens-before and lock-set analysis, or timing-based detection [20, 19, 28, 25, 13]. They could be used to approximate our reduction from data race checking to reachability in sequential programs. Some works [4, 22, 12] propose static analyses for finding data races. [4] designs a compositional data race detector for multi-threaded Java programs, based on an inter-procedural analysis assuming

that any two public methods can execute in parallel. Similar to [27], they precompute method summaries to extract potential racy accesses. These approaches are similar to the analysis in § 7, but they concern a different programming model.

**Analyzing Asynchronous Programs.** Several works propose program analyses for various classes of asynchronous programs. [7, 14] give complexity results for the reachability problem, and [27] proposes a static analysis for deadlock detection in C# programs that use both asynchronous and synchronous wait primitives. [8] investigates the problem of checking whether Java UI asynchronous programs have the same set of behaviors as sequential programs where roughly, asynchronous tasks are executed synchronously.

## 10 Conclusion

We proposed a framework for refactoring sequential programs to equivalent asynchronous programs based on async/await. We determined precise complexity bounds for the problem of computing all sound asynchronizations. This problem makes it possible to compute a sound asynchronization that maximizes performance by separating concerns – enumerate sound asynchronizations and evaluate performance separately. On the practical side, we have introduced an approximated synthesis procedure based on data-flow analysis that we implemented and evaluated on a benchmark of non-trivial C# programs.

The asynchronous programs rely exclusively on async/await and are deadlock-free by definition. Deadlocks can occur in a mix of async/await with "explicit" multi-threading that includes blocking `wait` primitives. Extending our approach for such programs is an interesting direction for future work.

## References

1. Bacon, D.F., Graham, S.L., Sharp, O.J.: Compiler transformations for high-performance computing. ACM Comput. Surv. **26**(4), 345–420 (1994). https://doi.org/10.1145/197405.197406, `https://doi.org/10.1145/197405.197406`
2. Beillahi, S.M., Bouajjani, A., Enea, C., Lahiri, S.: Artifact for the SAS 2022 paper: Automated Synthesis of Asynchronizations (May 2022). https://doi.org/10.5281/zenodo.7055422, `https://doi.org/10.5281/zenodo.7055422`
3. Bierman, G.M., Russo, C.V., Mainland, G., Meijer, E., Torgersen, M.: Pause 'n' play: Formalizing asynchronous c#. In: Noble, J. (ed.) ECOOP 2012 - Object-Oriented Programming - 26th European Conference, Beijing, China, June 11-16, 2012. Proceedings. Lecture Notes in Computer Science, vol. 7313, pp. 233–257. Springer (2012). https://doi.org/10.1007/978-3-642-31057-7_12, `https://doi.org/10.1007/978-3-642-31057-7\_12`
4. Blackshear, S., Gorogiannis, N., O'Hearn, P.W., Sergey, I.: Racerd: compositional static race detection. Proc. ACM Program. Lang. **2**(OOPSLA), 144:1–144:28 (2018). https://doi.org/10.1145/3276514, `https://doi.org/10.1145/3276514`
5. Bloem, R., Hofferek, G., Könighofer, B., Könighofer, R., Ausserlechner, S., Spork, R.: Synthesis of synchronization using uninterpreted functions. In: Formal Methods

in Computer-Aided Design, FMCAD 2014, Lausanne, Switzerland, October 21-24, 2014. pp. 35–42. IEEE (2014). https://doi.org/10.1109/FMCAD.2014.6987593, https://doi.org/10.1109/FMCAD.2014.6987593

6. Blume, W., Doallo, R., Eigenmann, R., Grout, J., Hoeflinger, J.P., Lawrence, T., Lee, J., Padua, D.A., Paek, Y., Pottenger, W.M., Rauchwerger, L., Tu, P.: Parallel programming with polaris. Computer **29**(12), 87–81 (1996). https://doi.org/10.1109/2.546612, https://doi.org/10.1109/2.546612
7. Bouajjani, A., Emmi, M.: Analysis of recursively parallel programs. In: Field, J., Hicks, M. (eds.) Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012. pp. 203–214. ACM (2012). https://doi.org/10.1145/2103656.2103681, https://doi.org/10.1145/2103656.2103681
8. Bouajjani, A., Emmi, M., Enea, C., Ozkan, B.K., Tasiran, S.: Verifying robustness of event-driven asynchronous programs against concurrency. In: Yang, H. (ed.) Programming Languages and Systems - 26th European Symposium on Programming, ESOP 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings. Lecture Notes in Computer Science, vol. 10201, pp. 170–200. Springer (2017). https://doi.org/10.1007/978-3-662-54434-1_7, https://doi.org/10.1007/978-3-662-54434-1_7
9. Cerný, P., Clarke, E.M., Henzinger, T.A., Radhakrishna, A., Ryzhyk, L., Samanta, R., Tarrach, T.: From non-preemptive to preemptive scheduling using synchronization synthesis. In: Kroening, D., Pasareanu, C.S. (eds.) Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part II. Lecture Notes in Computer Science, vol. 9207, pp. 180–197. Springer (2015). https://doi.org/10.1007/978-3-319-21668-3_11, https://doi.org/10.1007/978-3-319-21668-3_11
10. Cerný, P., Henzinger, T.A., Radhakrishna, A., Ryzhyk, L., Tarrach, T.: Regression-free synthesis for concurrency. In: Biere, A., Bloem, R. (eds.) Computer Aided Verification - 26th International Conference, CAV 2014, Held as Part of the Vienna Summer of Logic, VSL 2014, Vienna, Austria, July 18-22, 2014. Proceedings. Lecture Notes in Computer Science, vol. 8559, pp. 568–584. Springer (2014). https://doi.org/10.1007/978-3-319-08867-9_38, https://doi.org/10.1007/978-3-319-08867-9_38
11. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Grumberg, O., Veith, H. (eds.) 25 Years of Model Checking - History, Achievements, Perspectives. Lecture Notes in Computer Science, vol. 5000, pp. 196–215. Springer (2008). https://doi.org/10.1007/978-3-540-69850-0_12, https://doi.org/10.1007/978-3-540-69850-0_12
12. Engler, D.R., Ashcraft, K.: Racerx: effective, static detection of race conditions and deadlocks. In: Scott, M.L., Peterson, L.L. (eds.) Proceedings of the 19th ACM Symposium on Operating Systems Principles 2003, SOSP 2003, Bolton Landing, NY, USA, October 19-22, 2003. pp. 237–252. ACM (2003). https://doi.org/10.1145/945445.945468, https://doi.org/10.1145/945445.945468
13. Flanagan, C., Freund, S.N.: Fasttrack: efficient and precise dynamic race detection. In: Hind, M., Diwan, A. (eds.) Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2009, Dublin, Ireland, June 15-21, 2009. pp. 121–133. ACM (2009).

https://doi.org/10.1145/1542476.1542490, https://doi.org/10.1145/1542476.1542490

14. Ganty, P., Majumdar, R.: Algorithmic verification of asynchronous programs. ACM Trans. Program. Lang. Syst. **34**(1), 6:1–6:48 (2012). https://doi.org/10.1145/2160910.2160915, https://doi.org/10.1145/2160910.2160915
15. Godefroid, P., Yannakakis, M.: Analysis of boolean programs. In: Piterman, N., Smolka, S.A. (eds.) Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings. Lecture Notes in Computer Science, vol. 7795, pp. 214–229. Springer (2013). https://doi.org/10.1007/978-3-642-36742-7_16, https://doi.org/10.1007/978-3-642-36742-7_16
16. Gokhale, S., Turcotte, A., Tip, F.: Automatic migration from synchronous to asynchronous javascript apis. Proc. ACM Program. Lang. **5**(OOPSLA), 1–27 (2021). https://doi.org/10.1145/3485537, https://doi.org/10.1145/3485537
17. Gupta, A., Henzinger, T.A., Radhakrishna, A., Samanta, R., Tarrach, T.: Succinct representation of concurrent trace sets. In: Rajamani, S.K., Walker, D. (eds.) Proceedings of the 42nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2015, Mumbai, India, January 15-17, 2015. pp. 433–444. ACM (2015). https://doi.org/10.1145/2676726.2677008, https://doi.org/10.1145/2676726.2677008
18. Han, H., Tseng, C.: A comparison of parallelization techniques for irregular reductions. In: Proceedings of the 15th International Parallel & Distributed Processing Symposium (IPDPS-01), San Francisco, CA, USA, April 23-27, 2001. p. 27. IEEE Computer Society (2001). https://doi.org/10.1109/IPDPS.2001.924963, https://doi.org/10.1109/IPDPS.2001.924963
19. Kini, D., Mathur, U., Viswanathan, M.: Dynamic race prediction in linear time. In: Cohen, A., Vechev, M.T. (eds.) Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017. pp. 157–170. ACM (2017). https://doi.org/10.1145/3062341.3062374, https://doi.org/10.1145/3062341.3062374
20. Li, G., Lu, S., Musuvathi, M., Nath, S., Padhye, R.: Efficient scalable thread-safety-violation detection: finding thousands of concurrency bugs during testing. In: Brecht, T., Williamson, C. (eds.) Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019. pp. 162–180. ACM (2019). https://doi.org/10.1145/3341301.3359638, https://doi.org/10.1145/3341301.3359638
21. Lin, Y., Okur, S., Dig, D.: Study and refactoring of android asynchronous programming (T). In: Cohen, M.B., Grunske, L., Whalen, M. (eds.) 30th IEEE/ACM International Conference on Automated Software Engineering, ASE 2015, Lincoln, NE, USA, November 9-13, 2015. pp. 224–235. IEEE Computer Society (2015). https://doi.org/10.1109/ASE.2015.50, https://doi.org/10.1109/ASE.2015.50
22. Liu, B., Huang, J.: D4: fast concurrency debugging with parallel differential analysis. In: Foster, J.S., Grossman, D. (eds.) Proceedings of the 39th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2018, Philadelphia, PA, USA, June 18-22, 2018. pp. 359–373. ACM (2018). https://doi.org/10.1145/3192366.3192390, https://doi.org/10.1145/3192366.3192390

23. Manna, Z., Wolper, P.: Synthesis of communicating processes from temporal logic specifications. ACM Trans. Program. Lang. Syst. **6**(1), 68–93 (1984). https://doi.org/10.1145/357233.357237, https://doi.org/10.1145/357233.357237
24. Okur, S., Hartveld, D.L., Dig, D., van Deursen, A.: A study and toolkit for asynchronous programming in c#. In: Jalote, P., Briand, L.C., van der Hoek, A. (eds.) 36th International Conference on Software Engineering, ICSE '14, Hyderabad, India - May 31 - June 07, 2014. pp. 1117–1127. ACM (2014). https://doi.org/10.1145/2568225.2568309, https://doi.org/10.1145/2568225.2568309
25. Raman, R., Zhao, J., Sarkar, V., Vechev, M.T., Yahav, E.: Efficient data race detection for async-finish parallelism. In: Barringer, H., Falcone, Y., Finkbeiner, B., Havelund, K., Lee, I., Pace, G.J., Rosu, G., Sokolsky, O., Tillmann, N. (eds.) Runtime Verification - First International Conference, RV 2010, St. Julians, Malta, November 1-4, 2010. Proceedings. Lecture Notes in Computer Science, vol. 6418, pp. 368–383. Springer (2010). https://doi.org/10.1007/978-3-642-16612-9_28, https://doi.org/10.1007/978-3-642-16612-9_28
26. Roslyn: (2021), https://github.com/dotnet/roslyn
27. Santhiar, A., Kanade, A.: Static deadlock detection for asynchronous c# programs. In: Cohen, A., Vechev, M.T. (eds.) Proceedings of the 38th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2017, Barcelona, Spain, June 18-23, 2017. pp. 292–305. ACM (2017). https://doi.org/10.1145/3062341.3062361, https://doi.org/10.1145/3062341.3062361
28. Smaragdakis, Y., Evans, J., Sadowski, C., Yi, J., Flanagan, C.: Sound predictive race detection in polynomial time. In: Field, J., Hicks, M. (eds.) Proceedings of the 39th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2012, Philadelphia, Pennsylvania, USA, January 22-28, 2012. pp. 387–400. ACM (2012). https://doi.org/10.1145/2103656.2103702, https://doi.org/10.1145/2103656.2103702
29. Vechev, M.T., Yahav, E., Yorsh, G.: Inferring synchronization under limited observability. In: Kowalewski, S., Philippou, A. (eds.) Tools and Algorithms for the Construction and Analysis of Systems, 15th International Conference, TACAS 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings. Lecture Notes in Computer Science, vol. 5505, pp. 139–154. Springer (2009). https://doi.org/10.1007/978-3-642-00768-2_13, https://doi.org/10.1007/978-3-642-00768-2_13
30. Vechev, M.T., Yahav, E., Yorsh, G.: Abstraction-guided synthesis of synchronization. In: Hermenegildo, M.V., Palsberg, J. (eds.) Proceedings of the 37th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2010, Madrid, Spain, January 17-23, 2010. pp. 327–338. ACM (2010). https://doi.org/10.1145/1706299.1706338, https://doi.org/10.1145/1706299.1706338

## A Formalization and Proofs of Section 3

The following lemma shows that the absence of data races implies equivalence to the original program, in the sense of reaching the same set of configurations (program variable valuations).

**Lemma 2.** *$P_a[L_a]$ is sound implies $\mathbb{C}[P[L]] = \mathbb{C}[P_a[L_a]]$, for every $P_a[L_a] \in \mathsf{Asy}[P, L, L_a]$*

*Proof (Proof of Lemma 2).* Let $\rho$ be an execution of $P_a$ that reaches a configuration $\mathsf{ps} \in \mathbb{C}[P_a]$. We show that actions in $\rho$ can be reordered such that any action that occurs in $\rho$ between $(\_, i, \mathsf{call}(j))$ and $(\_, j, \mathsf{return})$ is not of the form $(\_, i, \_)$ (i.e., the task $j$ is executed synchronously). If an action $(\_, i, \_)$ occurs in $\rho$ between $(\_, i, \mathsf{call}(j))$ and $(\_, j, \mathsf{return})$, then it must be concurrent with $(j, \mathsf{return})$. Since $P_a$ does not admit data races, an execution $\rho'$ resulting from $\rho$ by reordering any two concurrent actions reaches the same configuration $\mathsf{ps}$ as $\rho$. Therefore, there exists an execution $\rho''$ where the actions that occur between any $(\_, i, \mathsf{call}(j))$ and $(\_, j, \mathsf{return})$ are not of the form $(\_, i, \_)$. This is also an execution of $P$ (modulo removing the awaits which have no effect), which implies $\mathsf{ps} \in \mathbb{C}[P]$.

## B Formalization and Proofs of Section 4

The following lemma shows that for a given $P_a$ there exists a unique $P'_a$ that is a maximal asynchronization of $P$ relative to $P_a$. The existence is implied by the fact that $\mathsf{strongAsy}[P, L, L_a]$ is the bottom element of $\leq$. To prove uniqueness, we assume by contradiction that there exist two incomparable maximal asynchronizations $P^1_a$ and $P^2_a$ and select the first await statement $s^1_w$, according to the control-flow of the sequential program, that is placed in different positions in the two programs. Assume that $s^1_w$ is closer to its matching call in $P^1_a$. Then, we move $s^1_w$ in $P^1_a$ further away from its matching call to the same position as in $P^2_a$. This modification does not introduce data races since $P^2_a$ is data race free. Thus, the resulting program is data race free, bigger than $P^1_a$, and smaller than $P_a$ w.r.t. $\leq$ contradicting the fact that $P^1_a$ is a maximal asynchronization.

**Lemma 3.** *Given an asynchronization $P_a \in \mathsf{Asy}[P, L, L_a]$, there exists a unique program $P'_a$ that is a maximal asynchronization of $P$ relative to $P_a$.*

*Proof (Proof of Lemma 3).* Since $\mathsf{strongAsy}[P, L, L_a]$ is the bottom element of $\leq$, then there always exists a sound asynchronization smaller than $P_a$. Assume by contradiction that there exist two distinct programs $P^1_a$ and $P^2_a$ that are both maximal asynchronizations of $P$ relative to $P_a$. Let $\rho^1$ (resp., $\rho^2$) be an execution of $P^1_a$ (resp., $P^2_a$) where every $\mathsf{await} * $ does not suspend the execution of the current task, i.e., $\rho^1$ and $\rho^2$ simulate the synchronous execution of $P$. Let $s^1_w$ be the statement corresponding to the first `await` action in $\rho^1$ such that (1) there exists an `await` action in $\rho^2$ with the corresponding `await` statement $s^2_w$, such that $s^1_w$ and $s^2_w$ match the same call in $P$, and $\mathsf{Cover}(s^1_w) \subset \mathsf{Cover}(s^2_w)$ (this

holds because $P_a^1$ and $P_a^2$ are distinct asynchronizations of the same synchronous program, thus $\mathsf{Cover}(s_w^1)$ and $\mathsf{Cover}(s_w^2)$ must be comparable), and (2) for every other `await` statement $s_w^3$ in $P_a^1$ that generates an `await` action which occurs before the `await` action of $s_w^1$ in $\rho^1$, there exists an `await` statement $s_w^4$ in $P_a^2$ matching the same call in $P$, such that $\mathsf{Cover}(s_w^3) = \mathsf{Cover}(s_w^4)$.

Let $P_a^3$ be the program obtained from $P_a^1$ by moving the await $s_w^1$ down (further away from the matching call) such that $\mathsf{Cover}(s_w^1) = \mathsf{Cover}(s_w^2)$. Moving an await down can only create data races between actions that occur after the execution of the matching call. Then, $P_a^3$ contains a data race iff there exists an execution $\rho$ of $P_a^3$ and two concurrent actions $a_1$ and $a_2$ that occur between the action $(\_, i, \mathsf{await}(j))$ generated by $s_w^1$ and the action $(\_, i, \mathsf{call}(j))$ of the call matching $s_w^1$, such that:

$$((\_, i, \mathsf{call}(j)), a_1) \in \mathsf{CO},\ (a_1, a_w) \notin \mathsf{HB},\ ((\_, i, \mathsf{call}(j)), a_2) \in \mathsf{CO} \text{ and } (a_2, (\_, i, \mathsf{await}(j))) \in \mathsf{HB}$$

where the action $a_w$ corresponds to the first await action in the task $j$. Let $s_w$ be the statement corresponding to the action $a_w$. Since the only difference between $P_a^3$ and $P_a^2$ is the placement of awaits then $((\_, i, \mathsf{call}(j)), a_1) \in \mathsf{CO}$ and $((\_, i, \mathsf{call}(j)), a_2) \in \mathsf{CO}$ hold in any execution $\rho'$ of $P_a^2$ that contains the actions $a_1$ and $a_2$. Also, note that since $a_w$ occurs in the task $j$ that the action of $s_w^1$ is waiting for. This implies that in $\rho^1$ the action of $s_w$ occurs before the action of $s_w^1$ in $\rho^1$. Therefore, by the definition of $s_w^1$ we have that $s_w$ in $P_a^1$ covers the same set of statements as the corresponding $s'_w$ in $P_a^2$ that matches the same call as $s_w$. Consequently, $(a_1, a'_w) \notin \mathsf{HB}$ and $(a_2, (\_, i, \mathsf{await}(j))) \in \mathsf{HB}$ hold in any execution $\rho'$ of $P_a^2$ that contains the actions $a_1$ and $a_2$ ($a'_w$ is the action of $s'_w$). Thus, there exists an execution $\rho'$ of $P_a^2$ such that the actions $a_1$ and $a_2$ are concurrent. This implies that if $P_a^3$ admits a data race, then $P_a^2$ admits a data race between actions generated by the same statements. As $P_a^2$ is data race free, we get that $P_a^3$ is data race free as well. Since $P_a^1 < P_a^3$, we get that $P_a^1$ is not maximal, which contradicts the hypothesis.

The complexity analysis also relies on a property of the maximal asynchronization relative to an immediate predecessor: if the predecessor is defined by moving an await $s''_w$, then the maximal asynchronization is obtained by moving only awaits smaller than $s''_w$ w.r.t. $\prec_w$.

**Lemma 4.** *If $P''_a$ is an immediate predecessor of a sound asynchronization $P'_a$, which is defined by moving an await $s''_w$ in $P'_a$ up, then the maximal sound asynchronization relative to $P''_a$ is obtained by moving only awaits smaller than $s''_w$ w.r.t. $\prec_w$.*

*Proof (Proof of Lemma 4).* Moving an await up in $P'_a$ can only create data races between actions that occur after the execution of this await (because the invocation is suspended earlier). The only possible repairs of these data races consists in either moving $s''_w$ down which results in $P'_a$ or moving up some other awaits that occur in methods that (indirectly) call the method in which $s''_w$ occurs. The first case is not applicable because it gives a program that is not

smaller than $P_a''$. In the second case, every await $s_w'$ that is moved up occurs in a method that (indirectly) calls the method in which $s_w''$ occurs, and therefore, $s_w'$ is smaller than $s_w''$ w.r.t. $\prec_w$.

Before giving the proof of Theorem 1, we note that the total order relation $\prec_w$ between awaits is fixed throughout the recursion of AsySyn and it corresponds to the order of the awaits in the weakest asynchronization of $P$, i.e., $\mathsf{wkAsy}[P, L, L_a]$. This is because the order between awaits in the same method might change from one asynchronization to another in $\mathsf{Asy}[P, L, L_a]$. If the control-flow graph of a method contains branches, it is possible to replace all `await` statements matching $s_c$ that are reachable in the CFG from $s$ with a single `await` statement $s_w$, in this case $s_w$ is ordered before any other await that one of the awaits that $s_w$ replaces is ordered before and is ordered after any await that all the awaits that $s_w$ replaces are ordered before. Also, it is possible to add additional `await`s statements in branches, in this case derive a total order between these awaits and order the awaits before or after any other await that the original await was ordered before or after, respectively.

*Proof (Proof of Theorem 1).* Let $P_a$ be the weakest asynchronization of $P$, then the set of all sound asynchronizations of $P$ is $\mathcal{A} = \{P_a'' : \mathbb{C}[P_a''] = \mathbb{C}[P] \text{ and } P_a'' \leq P_a'\}$, where $P_a'$ is the maximal asynchronization of $P$ relative to $P_a$. It is clear that every asynchronization outputted by AsySyn$(P_a, s_w^0)$ is in the set $\mathcal{A}$.

Let $P_a^0$ be a sound asynchronization of $P[L]$ w.r.t. $L_a$, i.e., $P_a^0 \in \mathcal{A}$. We will show that AsySyn outputs $P_a^0$. We have that either $P_a^0 = P_a'$ or $P_a^0 < P_a'$. The first case implies that $P_a^0$ is in AsySyn$(P_a, s_w)$. For the second case: let $s_w^{1'}$ be the maximum element in $P_a'$ w.r.t. $\prec_w$ that matches the same call as $s_w^{1'}$ in $P_a^0$ s.t. $\mathsf{Cover}(s_w^{1'}) \subset \mathsf{Cover}(s_w^1)$. Then, let $(P_a^1, s_w^1) \in \mathsf{ImPred}(P_a', s_w)$. We obtain that either $P_a^0 = P_a^1$ or $P_a^0 < P_a^1$. The fist case implies that $P_a^0$ is in AsySyn$(P_a, s_w)$. For the second case: let $P_a^{1'} = \textsc{MaxRel}(P_a^1)$ then either $P_a^0 = P_a^{1'}$ or $P_a^0 < P_a^{1'}$. The fist case implies that $P_a^0$ is in AsySyn$(P_a, s_w)$. For the second case: let $s_w^2$ be the maximum element in $P_a^{1'}$ w.r.t. $\prec_w$ that matches the same call as $s_w^{2'}$ in $P_a^0$ s.t. $\mathsf{Cover}(s_w^{2'}) \subset \mathsf{Cover}(s_w^2)$. Since $P_a^1$ is an immediate successor of $P_a'$ by moving the await $s_w^1$, then Lemma 4 implies $P_a^{1'}$ is obtained by moving only awaits smaller than $s_w^1$ w.r.t. $\prec_w$. Then, we either have $s_w^2 = s_w^1$ or $s_w^2 \prec_w s_w^1$. Thus, $(P_a^2, s_w^2) \in \mathsf{ImPred}(P_a^{1'}, s_w^1)$. We then obtain that either $P_a^0 = P_a^2$ or $P_a^0 < P_a^2$. Then, we repeat the above proof process until we obtain $P_a^n = P_a^0$. Thus, AsySyn outputs $P_a^0$.

Let $s_w^1$ and $s_w^2$ be two distinct await statements in $P_a'$ s.t. $s_w^2 \prec_w s_w^1$ and $(P_a^1, s_w^1)$, $(P_a^2, s_w^2) \in \mathsf{ImPred}(P_a', s_w)$. Similar to before then we have that $P_a^{2'} = \textsc{MaxRel}(P_a^2)$ is obtained by moving only awaits smaller than $s_w^2$ w.r.t. $\prec_w$. Thus, in $P_a^{2'}$ the await $s_w^1$ is in the same position as in $P_a'$. Then, $P_a^{1'} = \textsc{MaxRel}(P_a^1)$ is different than $P_a^{2'}$. For any two programs $P_a^{1''}$ and $P_a^{2''}$ s.t. $P_a^{1''}$ (resp., $P_a^{2''}$) is outputted by AsySyn$(P_a^{1'}, s_w^1)$ (resp., AsySyn$(P_a^{2'}, s_w^2)$), we have that the two programs are distinct since in $P_a^{2''}$ the await $s_w^1$ is in the

same position as in $P'_a$. Thus, we get that AsySyn outputs every element of $\mathcal{A}$ only once.

## C Formalization and Proofs of Section 5

The following lemma proves that for any unsound asynchronization, any trace with a data race contains at least one data race that involves two actions that are synchronously reachable. For instance, the program in Fig. 11 has two data races, one between $x = 1$ and $r4 = x$ and the other between $y = 2$ and $r5 = y$. However, the statement $y = 2$ is not reachable in the corresponding synchronous program. It is reachable in this asynchronization because of the data race between $x = 1$ and $r4 = x$, which are both reachable in the synchronous program. Eliminating the latter data race by moving the statement await $r1$ before $x = 1$, makes $y = 2$ unreachable and the data race between $y = 2$ and $r5 = y$ is also eliminated.

```
async method Main {
  r1 = call m;
  x = 1;
  await r1;
}
async method m {
  r2 = call m1;
  r3 = call m1;
  await r2;
  r4 = x;
  if r4 == 1
    y = 2;
  await r3;
}
async method m1 {
  await *;
  r5 = y;
  return;
}
```

Fig. 11

**Lemma 5.** *An asynchronization $P_a[L_a]$ is sound iff it does not admit data races between actions that are synchronously reachable.*

*Proof (Proof of Lemma 5).* Assume by contradiction that $P_a[L_a]$ is sound and it admits a data race $(a_1, a_2)$ in a trace $\tau \in \mathbb{T}\mathsf{r}(P_a[L_a])$ where one of the actions, say $a_1$, is not synchronously reachable. We assume w.l.o.g that the data race $(a_1, a_2)$ is the first that occurs in $\tau$ with at least one synchronously unreachable action. Then, there must exist a read access $a_r$ that enabled $a_1$, and therefore, $a_r$ reads a value that was not read in any synchronous execution. Thus, the read value must the result of another data race that occurs earlier in the trace $\tau$, which is a contradiction.

In Fig. 12, we explain how the repairing data races based on the partial order relating data races allows to avoid superfluous repair steps. For instance, in Fig. 12, the first data race to repair involves the read of `x` from `Main` and the write to `x` in `m`, because these statements are the first to execute in the original sequential program among the other statements involved in data races. Repairing this data race consists in moving `await r1` before the read of `x` from `Main`, which implies that `m` completes before the read of `x`. This repair is defined from a notion of *root cause* of a data race, that in this case, contains the call to `m` and the read of `x` from `Main`. Interestingly, this repair step removes the write-write data race between the write to `x` in `Main` and the write to `x` in `m` as well. If we would have repaired these data races in the opposite order, we would have moved `await t1` first before the write to `x`, and then, before the read of `x`.

```
async method Main {
 r1 = call m;
 r2 = x;
 x = r2 + 1;
 await r1;
}
async method m {
 await *;
 x = 2;
 return;
}
```

Fig. 12

In Fig. 13, we give a non-deterministic program where two statements of the program can be executed in different orders in different executions. In particular, the statements `r1 = x` and `r2 = y` of the program can be executed in different orders depending on the number of loop iterations and whether the if branch is entered during the first loop iteration.

```
method Main {
  while *
    if *
      r1 = x;
    r2 = y;
}
```

Fig. 13

For the program in Fig. 14, we have the following order between data races: $(\texttt{x = input}, \texttt{r2 = x}) \prec_{\mathsf{SO}} (\texttt{retVal = x}, \texttt{x = r2 + 1})$ because `r2 = x` is executed before the write `x = r2 + 1` in the original synchronous program (for simplicity we use statements instead of actions). However, the data races $(\texttt{x = input}, \texttt{r2 = x})$ and $(\texttt{x = input}, \texttt{r3 = x})$ are incomparable.

```
async method Main {
  r1 = call m;
  if *
    r2 = x;
    x = r2 + 1;
  else
    r3 = x;
  await r1;
}
async method m {
  await *
  retVal = x;
  x = input;
  return;
}
```

Fig. 14

The following lemma identifies a sufficient transformation for repairing a data race $(a_1, a_2)$: moving the await $s_w$ generating the action $a_w$ just before the statement $s$ generating $a$. This is sufficient because it ensures that every statement that follows $\mathsf{LCA}_{\mathsf{CO}}(a_1, a_2)$[10] in call order will be executed before $a$ and before any statement which succeeds $a$ in call order, including $a_2$. Note that moving the await $a_w$ anywhere after $a$ will not affect the concurrency between $a_1$ and $a_2$.

**Lemma 6.** *Let $(a_1, a_2)$ be a data race in a trace $\tau$ of an asynchronization $P_a$, and $a_c = (i, \mathsf{call}(j)) = \mathsf{LCA}_{\mathsf{CO}}(a_1, a_2)$. Then, $\tau$ contains a unique action $a_w = (i, \mathsf{await}(j))$ and a unique action $a$ such that:*

- *$(a, a_w) \in \mathsf{MO}$, and $a$ is the latest action in the method order $\mathsf{MO}$ such that $(a_c, a) \in \mathsf{MO}$ and $(a, a_2) \in \mathsf{CO}^*$ ($\mathsf{CO}^*$ denotes the reflexive closure of $\mathsf{CO}$).*

*Proof (Proof of Lemma 6).* Let $\rho$ be the execution of the trace $\tau$. By definition, $\rho$ ends with a configuration where the call stack and the set of pending tasks are empty. Therefore, $\rho$ contains an action $a_w = (\_, i, \mathsf{await}(j))$ matching $a_c$ which is unique by the definition of the semantics. Since $(a_c, a_1) \in \mathsf{CO}$ and $(a_c, a_2) \in \mathsf{CO}$ then either $a_c$ and $a_2$ occur in the same method, or there exists a call action $a'$ in the same task as $a_c$ such that $(a', a_2) \in \mathsf{CO}$. Then, we define $a = a_2$ in the first case, and $a$ as the latest action in the same task as $a_c$ such that $(a, a_2) \in \mathsf{CO}$ in the second case. We have that $(a, a_w) \in \mathsf{MO}$ because otherwise, $(a_w, a) \in \mathsf{MO}$ and $(a, a_2) \in \mathsf{CO}^*$ implies that $(a_1, a_2) \in \mathsf{HB}$ (because $(a_1, a_w) \in \mathsf{HB}$, and $\mathsf{MO}$ and $\mathsf{CO}$ are included in $\mathsf{HB}$), and this contradicts $a_1$ and $a_2$ being concurrent.

When the control-flow graph of the method contains branches, the construction of $\mathsf{RDR}(P_a, s_c, s)$ involves (1) replacing all `await` statements matching $s_c$ that are reachable in the CFG from $s$ with a single `await` statement placed just before $s$, and (2) adding additional `await` statements in branches that "conflict" with the branch containing $s$. This is to ensure the syntactic constraints described in Section 2. These additional `await` statements are at maximal distance from the corresponding call statement because of the maximality requirement.

[10] We abuse the terminology and make no distinction between statements and actions.

For instance, to repair the data race between `r2 = x` and `x = input` in the program on the left of Fig. 15, the statement `await r1` must be moved before `r2 = x` in the `if` branch, which implies that another await must be added on the `else` branch. The result is given on the right of Fig. 15.

```
async method Main {
  r1 = call m;
  if *
    r2 = x;
  else
    r3 = y;
  await r1;
}
async method m {
  await *
  retVal = x;
  x = input;
  return;
}
```

```
async method Main {
  r1 = call m;
  if *
    await r1;
    r2 = x;
  else
    r3 = y;
    await r1;
}
async method m {
  await *
  retVal = x;
  x = input;
  return;
}
```

Fig. 15: Examples of asynchronizations.

The following lemma shows that repairing a minimal data race cannot introduce smaller data races (w.r.t. $\prec_{\mathsf{SO}}$), which ensures some form of monotonicity when repairing minimal data races iteratively.

**Lemma 7.** *Let $P_a$ be an asynchronization, $(a_1, a_2)$ a data race in $P_a$ that is minimal w.r.t. $\prec_{\mathsf{SO}}$, and $(s_c, s)$ the root cause of $(a_1, a_2)$. Then, $\mathsf{RDR}(P_a, s_c, s)$ does not admit a data race that is smaller than $(a_1, a_2)$ w.r.t. $\prec_{\mathsf{SO}}$.*

Proof of Lemma 7. The only modification in the program $P'_a = \mathsf{RDR}(P_a, s_c, s)$ compared to $P_a$ is the movement of the await $s_w$ matching the call $s_c$ to be before the statement $s$ in a method $m$. The concurrency added in $P'_a$ that was not possible in $P_a$ is between actions $(a', a'')$ generated by statements $s'$ and $s''$, respectively, as shown in Fig. 16. W.l.o.g., we assume that $(a', a'') \in \mathsf{SO}$. The statements $s_1$ and $s_2$ are those generating $a_1$ and $a_2$, respectively. The statement $s'$ is related by $\mathsf{CO}^*$ to some statement in $m$ that follows $s$, and $s''$ is related by $\mathsf{CO}^*$ to some statement that follows the call to $m$ in the caller of $m$. Note that $s'$ is ordered by $\prec$ after $s_2$. Since $(a_1, a_2) \in \mathsf{SO}$ and $(a', a'') \in \mathsf{SO}$ then $s_2 \prec s''$ and $s_1 \prec s'$. Thus, any new data race $(a', a'')$ in $P'_a$ that was not reachable in $P_a$ is bigger than $(a_1, a_2)$. □

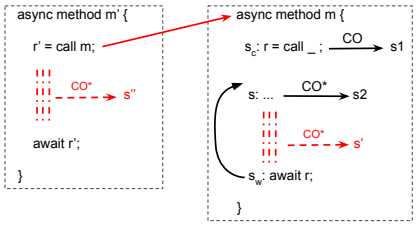


Fig. 16: An excerpt of an asynchronous program.

**Theorem 7.** *Given an asynchronization $P_a \in \mathsf{Asy}[P, L, L_a]$, MaxRel$(P_a)$ returns the optimal asynchronization of $P$ relative to $P_a$.*

*Proof (Proof of Theorem 7).* Since the recursive calls RCMinDR find all data races between synchronously reachable actions then the output $P'_a = \text{MaxRel}(P_a)$ is sound and therefore it is equivalent to $P$ (Lemma 2 and Lemma 5). Now we need to show that any successor $P^1_a$ of $P'_a$ that is also smaller than $P_a$ (w.r.t. $\leq$) admits data races. Let $s_w$ be the biggest await statement w.r.t. $\prec_w$ whose position in $P^1_a$ is changed with respect to its position in $P'_a$ (moved down). Since $P^1_a \leq P_a$, then $s_w$ was also moved up by the procedure MaxRel with respect to its position in $P_a$ to fix some data race $(a_1, a_2)$. Let $m$ be the method $m$ that

contains $s_w$ and $s_c$ be the matching call. We will now show that $(a_1, a_2)$ forms a data race in $P_a^1$ as well. $P_a^1$ has an execution $\rho$ that reaches both $a_1$ and $a_2$ (since $\mathbb{E}\mathsf{x}(P_a^1)$ includes the synchronous execution where all $\mathsf{await} * $ are interpreted as skip which reaches $a_1$ and $a_2$). Since every other await $s'_w$ in $P_a^1$ that occurs in a method $m'$ (in)directly called by $m$ (including the method associated with the call $s_c$) is in the same position as in $P'_a$, then the two actions $a_1$ and $a_2$ are not related by $\mathsf{HB}$ and are concurrent. Thus, $(a_1, a_2)$ forms a data race in $P_a^1$, which concludes the proof.

The fact that data races are enumerated in the order defined by $\prec_{\mathsf{SO}}$ guarantees a bound on the number of times an await matching the same call is moved during the execution of MaxRel$(P_a)$. In general, this bound is the number of statements covered by all the awaits matching the call in the input program $P_a$. Actually, this is a rather coarse bound. A more refined analysis has to take into account the number of branches in the CFGs. For programs without conditionals or loops, every await is moved at most once during the execution of MaxRel$(P_a)$. In the presence of branches, a call to an asynchronous method may match multiple `await` statements (one for each CFG path starting from the call), and the data races that these `await` statements may create may be incomparable w.r.t. $\prec_{\mathsf{SO}}$. Therefore, for a call statement $s_c$, let $|s_c|$ be the sum of $|\mathsf{Cover}(s_w)|$ for every await $s_w$ matching $s_c$ in $P_a$.

**Lemma 8.** *For any asynchronization $P_a \in \mathsf{Asy}[P, L, L_a]$ and call statement $s_c$ in $P_a$, the while loop in* MaxRel$(P_a)$ *does at most $|s_c|$ iterations that result in moving an await matching $s_c$.*

*Proof (Proof of Lemma 8).* We consider first the case without conditionals or loops, and we show by contradiction that every `await` statement $s_w$ is moved at most once during the execution of MaxRel$(P_a)$, i.e., there exists at most one iteration of the while loop which changes the position of $s_w$. Suppose that the contrary holds for an `await` $s_w$. Let $(a_1, a_2)$, and $(a_3, a_4)$ be the data races repaired by the first and second moves of $s_w$, respectively. By Lemma 6, there exist two actions $a$ and $a'$ such that

$$(a_c, a) \in \mathsf{MO},\ (a, a_2) \in \mathsf{CO}^*, (a, a_w) \in \mathsf{MO} \text{ and } (a_c, a') \in \mathsf{MO},\ (a', a_4) \in \mathsf{CO}^*, (a', a_w) \in \mathsf{MO}$$

where $a_w = (\_, i, \mathsf{await}(j))$ and $a_c = (\_, i, \mathsf{call}(j))$ are the asynchronous call action and the matching await action. Let $s_2$ and $s_4$ be the statements generating the two actions $a_2$ and $a_4$, respectively. Then, we have either $s_2 \prec s_4$ or $s_2 = s_4$, and both cases imply that $(a, a') \in \mathsf{MO}^*$. Thus, moving the await statement generating $a_w$ before the statement generating $a$ implies that it is also placed before the statement generating $a'$ (that occurs after $a$ in the same method). Thus, the first move of the await $s_w$ repaired both data races, which is contradiction.

In the presence of conditionals or loops, moving an await up in one branch may correspond to adding multiple awaits in the other conflicting branches. Also, one call in the program may correspond to multiple awaits on different branches. However, every repair of a data race consists in moving one await closer to the matching call $s_c$ and before one more statement covered by some await matching $s_c$ in the input $P_a$.

```
1  Add before s1:
2  if ( lastTaskDelayed == ⊥ && * )
3    lastTaskDelayed := myTaskId();
4    DescendantDidAwait := thisHasDoneAwait;
5    return

7  Add before s2:
8    if ( task_sc == myTaskId() )
9      s := s2;
10   assert (lastTaskDelayed == ⊥ ||
         !DescendantDidAwait);

13 Replace every statement ‘‘await r’’ with:
14   if( r == lastTaskDelayed  ) then
15     if ( !DescendantDidAwait )
16       DescendantDidAwait :=
        thisHasDoneAwait;
17     lastTaskDelayed := myTaskId();
18     return
19   else
20     thisHasDoneAwait := true

22 Add before every statement ‘‘r := call m’’:
23   if ( task_sc == myTaskId() ) then
24     s := this statement;

26 Add after every statement ‘‘r := call m’’:
27   if ( r == lastTaskDelayed )
28     sc := this statement;
29     task_sc := myTaskId();
```

Fig. 17: A program instrumentation for computing the root cause of a minimal data race between the statements $s_1$ and $s_2$ (if any). All variables except for `thisHasDoneAwait` are program (global) variables. `thisHasDoneAwait` is a local variable. The value $\bot$ represents an initial value of a variable. The variables $s_c$ and $s$ store the (program counters of the) statements representing the root cause. The method `myTaskId` returns the id of the current task.

## D Computing Root Causes of Minimal Data Races

We present a reduction from the problem of computing root causes of minimal data races to reachability (assertion checking) in sequential programs. This reduction builds on a program instrumentation for checking if there exists a minimal data race that involves two given statements $(s_1, s_2)$ that are reachable in an execution of the original synchronous program, whose correctness relies on the assumption that another pair of statements cannot produce a smaller data race. This instrumentation is used in an iterative process where pairs of statements are enumerated according to the colexicographic order induced by $\prec$. This specific enumeration ensures that the assumption made for the correctness of the instrumentation is satisfied.

Given an asynchronization $P_a$, the instrumentation described in Fig. 17 represents a synchronous program where all `await` statements are replaced with synchronous code (lines 14–20). This instrumentation simulates asynchronous executions of $P_a$ where methods may be only partially executed, modeling `await` interruptions. It reaches an error state (see the assert at line 10) when an action generated by $s_1$ is concurrent with an action generated by $s_2$, which represents a data race, provided that $s_1$ and $s_2$ access a common program variable (these statements are assumed to be given as input). Also, the values of $s_c$ and $s$ when reaching the assertion violation represent the root-cause of this data race.

The instrumentation simulates an execution of $P_a$ to search for a data race as follows (we discuss the identification of the root-cause afterwards):

– It executes under the synchronous semantics until an instance of $s_1$ is nondeterministically chosen as a candidate for the first action in the data race ($s_1$ can execute multiple times if it is included in a loop for instance). The

current invocation is interrupted when it is about to execute this instance of $s_1$ and its task id $t_0$ is stored into `lastTaskDelayed` (see lines 2–5).
- Every invocation that transitively called $t_0$ is interrupted when an `await` for an invocation in this call chain (whose task id is stored into `lastTaskDelayed`) would have been executed in the asynchronization $P_a$ (see line 18).
- Every other method invocation is executed until completion as in the synchronous semantics.
- When reaching $s_2$, if $s_1$ has already been executed (`lastTaskDelayed` is not $\perp$) and at least one invocation has only partially been executed, which is recorded in the boolean flag `DescendantDidAwait` and which means that $s_1$ is concurrent with $s_2$, then the instrumentation stops with an assertion violation.

A subtle point is that the instrumentation may execute code that follows an `await` $r$ even if the task $r$ has been executed only partially, which would not happen in an execution of the original $P_a$. Here, we rely on the assumption that there exist no data race between that code and the rest of the task $r$. Such data races would necessarily involve two statements which are before $s_2$ w.r.t. $\prec$. Therefore, the instrumentation is correct only if it is applied by enumerating pairs of statements $(s_1, s_2)$ w.r.t. the colexicographic order induced by $\prec$.

Next, we describe the computation of the root-cause, i.e., the updates on the variables $s_c$ and $s$. By definition, the statement $s_c$ in the root-cause should be a call that makes an invocation that is in the call stack when $s_1$ is reached. This can be checked using the variable `lastTaskDelayed` that stores the id of the last such invocation popped from the call stack (see the test at line 27). The statement $s$ in the root-cause can be any call statement that has been executed in the same task as $s_c$ (see the test at line 23), or $s_2$ itself (see line 9).

Let $[\![P_a, s_1, s_2]\!]$ denote the instrumentation in Fig. 17. We say that the values of $s_c$ and $s$ when reaching the assertion violation are the root cause computed by this instrumentation. The following theorem states its correctness.

**Theorem 8.** *If $[\![P_a, s_1, s_2]\!]$ reaches an assertion violation, then it computes the root cause of a minimal data race, or there exists $(s_3, s_4)$ such that $[\![P_a, s_3, s_4]\!]$ reaches an assertion violation and $(s_3, s_4)$ is before $(s_1, s_2)$ in colexicographic order w.r.t. $\prec$.*

Based on Theorem 8, we define an implementation of the procedure RCMinDR($P_a$) used in computing maximal asynchronizations (Algorithm 2) as follows:

- For all pairs of read or write statements $(s_1, s_2)$ in colexicographic order w.r.t. $\prec$ that are reachable in an execution of the original synchronous program $P$.
  - If $[\![P_a, s_1, s_2]\!]$ reaches an assertion violation, then
    - return the root cause computed by $[\![P_a, s_1, s_2]\!]$
- return $\perp$

Checking whether read or write statements are reachable can be determined using a linear number of reachability queries in the synchronous program $P$. Also,

the order $\prec$ between read or write statements can be computed using a quadratic number of reachability queries in the synchronous program $P$. Therefore, $s \prec s'$ iff an instrumentation of $P$ that sets a flag when executing $s$ and asserts that this flag is not set when executing $s'$ reaches an assertion violation. The following theorem states the correctness of the procedure above.

**Theorem 9.** RCMinDR($P_a$) *returns the root cause of a minimal data race of* $P_a$ *w.r.t.* $\prec_{\mathsf{SO}}$*, or* $\bot$ *if* $P'_a$ *is data race free.*

## E Formalization and Proofs of Section 6

**Theorem 10.** *Checking whether there exists a sound asynchronization different from the strong asynchronization is PSPACE-complete.*

*Proof (Proof of Theorem 10).* (1) define a new method $m$ that writes to a new program variable $x$, and insert a call to $m$ followed by a write to $x$ at location $\ell$, and (2) insert a write to $x$ after every call statement that calls a method in $\{m'\}^*$, where $m'$ is the method containing $\ell$. Let $m_a$ be an asynchronous version of $m$ obtained by inserting an `await` $*$ at the beginning. Then, $\ell$ is reachable in $P$ iff the only sound asynchronization of $P'$ w.r.t. $\{m_a\}$ is the strong asynchronization.

## F Formalization and Proofs of Section 7

The MaxRel$^{\#}$ procedure repairs data races in an order which is $\prec_{\mathsf{SO}}$ with some exceptions that do not affect optimality, i.e., the number of times an await matching the same call can be moved. For instance, if a method $m$ calls two other methods $m_1$ and $m_2$ in this order, the procedure above may handle $m_2$ before $m_1$, i.e., repair data races between actions that originate from $m_2$ before data races that originate from $m_1$, although the former are bigger than the latter in $\prec_{\mathsf{SO}}$. This does not affect optimality because those repairs are "independent", i.e., any repair in $m_2$ cannot influence a repair in $m_1$, and vice-versa. The crucial point is that this procedure repairs data races between actions that originate from a method $m$ before data races that involve actions in methods preceding $m$ in the call graph, which are bigger in $\prec_{\mathsf{SO}}$ than the former.

Note that MaxRel$^{\#}$ procedure which is based on the bottom-up inter-procedural data-flow analysis compromises precision to reduce the complexity of the problem from undecidable in general or PSPACE-complete with finite data to polynomial time. However, because of this imprecision, certain await statements may be moved closer to the matching call unnecessarily. For instance, in Fig. 11, the precise algorithm (using the procedure MaxRel in Algorithm 2) will only repair the data race on $x$ because doing so, the potential data race on $y$ will become unreachable. On the other hand, the polynomial-time algorithm (using the MaxRel$^{\#}$ procedure) will also repair the data race on $y$, moving another await closer to the matching call, since it cannot reason about data (one statement of this data race is only reachable if the variable $r4$ is 1).