

# Fence Synthesis under the C11 Memory Model

Sanjana Singh<sup>1</sup>, Divyanjali Sharma<sup>1</sup> Ishita Jaju<sup>2</sup>, and Subodh Sharma<sup>1</sup>

<sup>1</sup> Indian Institute of Technology Delhi, India,  
{sanjana.singh,divyanjali,svs}@cse.iitd.ac.in,  
<sup>2</sup> Uppsala University, Sweden

**Abstract.** The C/C++11 (*C11*) standard offers a spectrum of ordering guarantees on memory access operations. The combinations of such orderings pose a challenge in developing *correct* and *efficient* weak memory programs. A common solution to preclude those program outcomes that violate the correctness specification is using *C11* synchronization-fences, which establish ordering on program events. The challenge is in choosing a combination of fences that (i) restores the correctness of the input program, with (ii) as little impact on efficiency as possible (*i.e.*, the smallest set of weakest fences). This problem is the *optimal fence synthesis* problem and is NP-hard for straight-line programs. In this work, we propose the first fence synthesis technique for *C11* programs called **FenSyng** and show its optimality. We additionally propose a near-optimal efficient alternative called **fFenSyng**. We prove the optimality of **FenSyng** and the soundness of **fFenSyng** and present an implementation of both techniques. Finally, we contrast the performance of the two techniques and empirically demonstrate **fFenSyng**'s effectiveness.

**Keywords:** *C11*, fence-synthesis, optimal

## 1 Introduction

Developing weak memory programs requires careful placement of fences and memory barriers to preserve ordering between program instructions and exclude undesirable program outcomes. However, computing the correct combination of the type and location of fences is challenging. Too few or incorrectly placed fences may not preserve the necessary ordering, while too many fences can negatively impact the performance. Striking a balance between preserving the correctness and obtaining performance is highly non-trivial even for expert programmers.

This paper presents an automated fence synthesis solution for weak memory programs developed using the C/C++11 standard (*C11*). *C11* provides a spectrum of ordering guarantees called *memory orders*. In a program, a memory access operation is associated with a memory order which specifies how other memory accesses are ordered with respect to the operation. The memory orders range from *relaxed* (**rlx**) (that imposes no ordering restriction) to *sequentially-consistent* (**sc**) (that may restore sequential consistency). Understanding all the subtle complexities of *C11* orderings and predicting the program outcomes can quickly become exacting. Consider the program (**RWRW**) (§2), where the orders

are shown as subscripts. When all the memory accesses are ordered `rlx`, there exists a program outcome that violates the correctness specification (specified as an *assert* statement). However, when all accesses are ordered `sc`, the program is provably correct.

In addition, the *C11* memory model supports *C11 fences* that serve as tools for imposing ordering restrictions. Notably, *C11* associates fences with memory orders, thus, supporting various degrees of ordering guarantees through fences.

This work proposes an *optimal* fence synthesis technique for *C11* called **FenSyng**. It involves finding solutions to two problems: (i) computing an optimal (minimal) set of locations to synthesize fences and (ii) computing an optimal (weakest) memory order to be associated with the fences (formally defined in §3). **FenSyng** takes as input *all* program runs that violate user-specified assertions and attempts optimal *C11* fence synthesis to stop the violating outcomes. **FenSyng** reports when *C11* fences alone cannot fix a violation. In general, computing a minimal number of fences with multiple types of fences is shown to be NP-hard for straight-line programs [24]. We note, rather unsurprisingly, that this hardness manifests in the proposed optimal fence synthesis solution even for the simplest *C11* programs. Our experiments (§7) show an exponential increase in the analysis time with the increase in the program size.

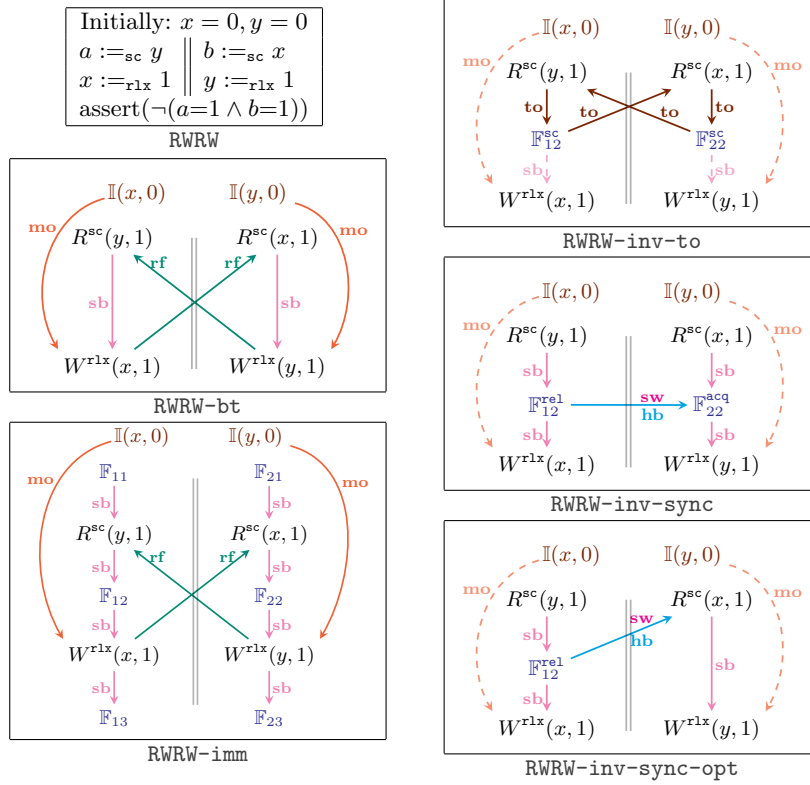
Further, to address scalability, this paper proposes a *near-optimal* fence synthesis technique called **fFenSyng** (fast **FenSyng**) that fixes *one* violating outcome at a time optimally. Note that fixing one outcome optimally may not guarantee optimality across all violating outcomes. In the process, this technique may add a small number of extra fences than what an optimal solution would compute. Our experiments reveal that **fFenSyng** performs exponentially better than **FenSyng** in terms of the analysis time while adding no extra fences in over 99.5% of the experiments.

Both **FenSyng** and **fFenSyng**, compute the solution from a set of combinations of fences that can stop the violating outcomes, also called *candidate solutions*. The candidate solutions are encoded in a *head-cycle-free* CNF SAT query [8]. Computing an optimal solution from candidates then becomes finding a solution to a *min-model finding problem*.

Many prior works have focused on automating fence synthesis (discussed in §8). However, the techniques presented in this paper are distinct from prior works in the following two ways: (i) prior techniques do not support *C11* memory orders, and (ii) the proposed techniques in this paper synthesize fences that are portable and not architecture-specific.

**Contributions.** To summarize, this work makes the following contributions:

- The paper presents **FenSyng** and **fFenSyng** (§6). To the best of our knowledge, these are the first fence synthesis techniques for *C11*.
- The paper shows (using Theorems 1 and 2) that the techniques are sound, *i.e.*, if the input program can be fixed by *C11* fences, then the techniques will indeed find a solution. The paper also shows (using Theorem 3) that **FenSyng** produces an optimal result in the number and type of fences.



- Finally, the paper presents an implementation of the said techniques and presents an empirical validation using a set of 1389 litmus tests. Further, the paper empirically shows the effectiveness of **fFenSyng** on a set of challenging benchmarks from prior works. **fFenSyng** performs on an average 67x faster than **FenSyng**.

## 2 Overview of FenSyng and fFenSyng

Given a program  $P$ , a *trace*  $\tau$  of  $P$  (formally defined in §3); is considered *buggy* if it violates an assertion of  $P$ . **FenSyng** takes *all* buggy traces of  $P$  as input. The difference in **fFenSyng** is that the input is a *single* buggy trace of  $P$ .

Consider the input program (**RWRW**), where  $x$  and  $y$  are shared objects with initial values 0, and  $a$  and  $b$  are local objects. Let  $W^m(o, v)$  and  $R^m(o, v)$  represent the write and read of object  $o$  and value  $v$  with the memory order  $m$ . Let  $I(o, v)$  represent the initialization event for object  $o$  with value  $v$ . The parallel bars ( $\parallel$ ) represent the parallel composition of events from separate threads. Figure (**RWRW-bt**) represents a buggy trace  $\tau$  of (**RWRW**) under *C11* semantics. For convenience, the relations – *sequenced-before* ( $\rightarrow_{\tau}^{sb}$ ), *reads-from* ( $\rightarrow_{\tau}^{rf}$ ), *modification-order* ( $\rightarrow_{\tau}^{mo}$ ) (formally defined in §3, §4) – among the events of  $\tau$  are shown. The

assert condition of (RWRW) is violated as the read events are not ordered before the write events of the same object, allowing reads from *later* writes.

Consider the following three sets of fences that can invalidate the trace (RWRW-bt):  $c_1 = \{\mathbb{F}_{12}^{\text{sc}}, \mathbb{F}_{22}^{\text{sc}}\}$ ,  $c_2 = \{\mathbb{F}_{12}^{\text{rel}}, \mathbb{F}_{22}^{\text{acq}}\}$  and  $c_3 = \{\mathbb{F}_{12}^{\text{rel}}\}$  (the superscripts indicate the memory orders and the subscripts represent the synthesis locations for fences). The solutions are depicted in Figures (RWRW-inv-to), (RWRW-inv-sync) and (RWRW-inv-sync-opt) for  $c_1$ ,  $c_2$  and  $c_3$ , respectively. The candidate solution  $c_1$  prevents a total order on the **sc** ordered events [9,12], thus, invalidating (RWRW-inv-to) under *C11* semantics. With candidate solution  $c_2$ , a happens-before ( $\rightarrow_{\tau}^{\text{hb}}$ ) ordering is formed (refer to §4) between  $R^{\text{sc}}(y, 1)$  and  $W^{\text{sc}}(y, 1)$ . This forbids a read from an ordered-later write, thus, invalidating (RWRW-inv-sync). Candidate solution  $c_3$  establishes a similar  $\rightarrow_{\tau}^{\text{hb}}$  ordering by exploiting the strong memory order of  $R^{\text{sc}}(x, 1)$  and invalidates (RWRW-inv-sync-opt).

The candidate solution  $c_2$  is preferred over  $c_1$  as it contains weaker fences. On the other hand, candidate  $c_3$  represents an optimal solution as it uses the smallest number of weakest fences. We formally define the optimality of fence synthesis in §3. While **FenSyng** will compute the solution  $c_3$ , **fFenSyng** may compute one from the many candidate solutions.

Both **FenSyng** and **fFenSyng** start by transforming each buggy trace  $\tau$  to an *intermediate* version,  $\tau^{\text{imm}}$ , by inserting *untyped C11 fences* (called *candidate fences*) above and below the trace events. (RWRW-imm) shows an intermediate version corresponding to (RWRW-bt). The addition of fences (assuming they are of the strongest variety) leads to the creation of new  $\rightarrow_{\tau}^{\text{hb}}$  ordering edges. This may result in cycles in the dependency graph under the *C11* semantics (explained in §4). The set of fences in a cycle constitutes a *candidate solution*. For example, an ordering from  $\mathbb{F}_{12}$  to  $\mathbb{F}_{22}$  in (RWRW-imm) induces a cyclic relation  $W^{\text{rlx}}(y, 1) \rightarrow_{\tau}^{\text{rf}} R^{\text{sc}}(y, 1) \rightarrow_{\tau}^{\text{hb}} W^{\text{rlx}}(y, 1)$  violating the  $\rightarrow_{\tau}^{\text{rf}}; \rightarrow_{\tau}^{\text{hb}}$  irreflexivity (refer to §4).

The candidate solutions are collected in a SAT query ( $\Phi$ ). Assuming  $c_1$ ,  $c_2$  and  $c_3$  are the only candidate solutions for (RWRW-bt), then  $\Phi = (\mathbb{F}_{12} \wedge \mathbb{F}_{22}) \vee (\mathbb{F}_{12} \wedge \mathbb{F}_{22}) \vee (\mathbb{F}_{12})$ , where for a fence  $\mathbb{F}_i^m$ ,  $\mathbb{F}_i$  represents the same fence with unassigned memory order. **fFenSyng** uses a SAT solver to compute the *min-model* of  $\Phi$ ,  $\text{min}\Phi = \{\mathbb{F}_{12}\}$ . Further, **fFenSyng** applies the *C11* ordering rules on fences to determine the weakest memory order for the fences in  $\text{min}\Phi$ . For instance,  $\mathbb{F}_{12}$  in  $\text{min}\Phi$  is computed to have the order **rel** (explained in §6). **fFenSyng** then inserts  $\mathbb{F}_{12}$  with memory order **rel** in (RWRW) at the location depicted in (RWRW-inv-sync-opt). This process repeats for the next buggy trace.

In contrast, since **FenSyng** works with all buggy traces at once, it requires the conjunction of the SAT queries  $\Phi_i$  corresponding to each buggy trace  $\tau_i$ . The min-model of the conjunction is computed, which provides optimality.

### 3 Preliminaries

Consider a multi-threaded *C11* program ( $P$ ). Each thread of  $P$  performs a sequence of *events* that are runtime instances of memory access operations (reads, writes, and rmws) on shared objects and *C11* fences. Note that an event is

uniquely identified in a trace; however, multiple events may be associated with the same program location. The events may be atomic or non-atomic.

**C11 memory orders.** The atomic events and fence operations are associated with memory orders that define the ordering restriction on atomic and non-atomic events around them. Let  $\mathcal{M} = \{\mathbf{na}, \mathbf{rlx}, \mathbf{rel}, \mathbf{acq}, \mathbf{ar}, \mathbf{sc}\}$ , represent the orders relaxed ( $\mathbf{rlx}$ ), release ( $\mathbf{rel}$ ), acquire/consume ( $\mathbf{acq}$ ), acquire-release ( $\mathbf{ar}$ ) and sequentially consistent ( $\mathbf{sc}$ ) for atomic events. A non-atomic event is recognized by the  $\mathbf{na}$  memory order. Let  $\sqsubset \subseteq \mathcal{M} \times \mathcal{M}$  represent the relation *weaker* such that  $m_1 \sqsubset m_2$  represents that the  $m_1$  is weaker than  $m_2$ . As a consequence, annotating an event with  $m_2$  may order two events that remain unordered with  $m_1$ . The orders in  $\mathcal{M}$  are related as  $\mathbf{na} \sqsubset \mathbf{rlx} \sqsubset \{\mathbf{rel}, \mathbf{acq}\} \sqsubset \mathbf{ar} \sqsubset \mathbf{sc}$ . We also define the relation  $\sqsubseteq$  to represent *weaker or equally weak*. Similarly, we define  $\sqsupset$  to represent *stronger* and  $\sqsupseteq$  to represent *stronger or equally strong*.

We use  $\mathcal{E}^{\mathbb{W}} \subseteq \mathcal{E}$  to denote the set of events that perform write to shared memory objects *i.e.*, write events or rmw events. Similarly, we use  $\mathcal{E}^{\mathbb{R}} \subseteq \mathcal{E}$  to denote events that read from a shared memory object *i.e.*, read events and rmw events, and  $\mathcal{E}^{\mathbb{F}}$  to denote the fence events. We also use  $\mathcal{E}^{(m)} \in \mathcal{E}$  (and accordingly  $\mathcal{E}^{\mathbb{W}(m)}$ ,  $\mathcal{E}^{\mathbb{R}(m)}$  and  $\mathcal{E}^{\mathbb{F}(m)}$ ) to represent the events with the memory order  $m \in \mathcal{M}$ ; as an example  $\mathcal{E}^{\mathbb{F}(\mathbf{sc})}$  represents the set of fences with the memory order  $\mathbf{sc}$ .

**Definition 1 (Trace).** A trace,  $\tau$ , of  $P$  is a tuple  $\langle \mathcal{E}_\tau, \rightarrow_\tau^{\mathbf{hb}}, \rightarrow_\tau^{\mathbf{mo}}, \rightarrow_\tau^{\mathbf{rf}} \rangle$ , where  $\mathcal{E}_\tau \subseteq \mathcal{E}$  represents the set of events in the trace  $\tau$ ;

- $\rightarrow_\tau^{\mathbf{hb}}$  (*Happens-before*)  $\subseteq \mathcal{E}_\tau \times \mathcal{E}_\tau$  is a partial order which captures the event interactions and inter-thread synchronizations, discussed in §4;
- $\rightarrow_\tau^{\mathbf{mo}}$  (*Modification-order*)  $\subseteq \mathcal{E}_\tau^{\mathbb{W}} \times \mathcal{E}_\tau^{\mathbb{W}}$  is a total order on the writes of an object;
- $\rightarrow_\tau^{\mathbf{rf}}$  (*Reads-from*)  $\subseteq \mathcal{E}_\tau^{\mathbb{W}} \times \mathcal{E}_\tau^{\mathbb{R}}$  is a relation from a write event to a read event signifying that the read event takes its value from the write event in  $\tau$ .

Note that, we use  $\mathcal{E}_\tau^{\mathbb{W}}$ ,  $\mathcal{E}_\tau^{\mathbb{R}}$  and  $\mathcal{E}_\tau^{\mathbb{F}}$  (and also  $\mathcal{E}_\tau^{\mathbb{W}(m)}$ ,  $\mathcal{E}_\tau^{\mathbb{R}(m)}$  and  $\mathcal{E}_\tau^{\mathbb{F}(m)}$  where  $m \in \mathcal{M}$ ) for the respective sets of events for a trace  $\tau$ .

**Relational Operators.**  $R^{-1}$  represents the inverse and  $R^+$  represents the transitive closure of a relation  $R$ . Further,  $R_1; R_2$  represents the composition of relations  $R_1$  and  $R_2$ . Let  $R|_{\mathbf{sc}}$  represent a subset of a relation  $R$  on  $\mathbf{sc}$  ordered events; *i.e.*  $(e_1, e_2) \in R|_{\mathbf{sc}} \iff (e_1, e_2) \in R \wedge e_1, e_2 \in \mathcal{E}^{(\mathbf{sc})}$ . Note that we also use the infix notation  $e_1 R e_2$  for  $(e_1, e_2) \in R$ . Lastly, a relation  $R$  has a cycle (or is cyclic) if  $\exists e_1, e_2 \in \mathcal{E}$  s.t.  $e_1 R e_2 \wedge e_2 R e_1$ .

**A note on optimality.** The notion of optimality may vary with context. Consider two candidate solutions  $\{\mathbb{F}_i^{\mathbf{sc}}\}$  and  $\{\mathbb{F}_j^{\mathbf{rel}}, \mathbb{F}_k^{\mathbf{acq}}\}$  where the superscripts represent the memory orders. The two solutions are incomparable under *C11*, and their performance efficiency is subject to the input program and the underlying architecture. **FenSyng** chooses a candidate solution  $c$  as an optimal solution if: (i)  $c$  has the smallest number of candidate fences, and (ii) each fence of  $c$  has the weakest memory order compared to other candidate solutions that satisfy (i).

Let  $sz(c)$  represent the size of the candidate solution  $c$  and given the set of all candidate solutions  $\{c_1, \dots, c_n\}$  to fix  $P$ , let  $sz(P) = \min(sz(c_1), \dots, sz(c_n))$ . Further, we assign weights  $wt(c)$  to each candidate solution  $c$ , computed as the summation of the weights of its fences where a fence ordered  $\mathbf{rel}$  or  $\mathbf{acq}$  is

assigned the weight 1, a fence ordered **ar** is assigned 2, and a fence ordered **sc** is assigned 3. Optimality for **FenSyng** is formally defined as:

**Definition 2. Optimality of fence synthesis.** Consider a set of candidate solutions  $c_1, \dots, c_n$ . A solution  $c_i$  (for  $i \in [1, n]$ ) is considered optimal if:

- (i)  $sz(c_i) = \underline{sz}(P) \wedge$  (ii)  $\forall j \in [1, n]$  s.t.  $sz(c_j) = \underline{sz}(P)$ ,  $wt(c_i) \leq wt(c_j)$ .

## 4 Background: C11 Memory Model

The *C11* memory model defines a trace using a set of event relations, described in Definition 1. The most significant relation that defines a *C11* trace  $\tau$  is the irreflexive and acyclic happens-before relation,  $\rightarrow_{\tau}^{\text{hb}} \subseteq \mathcal{E}_{\tau} \times \mathcal{E}_{\tau}$ . The  $\rightarrow_{\tau}^{\text{hb}}$  relation is composed of the following relations [12].

- $\rightarrow_{\tau}^{\text{sb}}$  (*Sequenced-before*): total occurrence order on the events of a thread.
- $\rightarrow_{\tau}^{\text{sw}}$  (*Synchronizes-with*) Inter-thread synchronization between a write  $e_w$  (ordered  $\sqsupseteq$  **rel**) and a read  $e_r$  (ordered  $\sqsupseteq$  **acq**) when  $e_w \rightarrow_{\tau}^{\text{rf}} e_r$ .
- $\rightarrow_{\tau}^{\text{dob}}$  (*Dependency-ordered-before*): Inter-thread synchronization between a write  $e_w$  (ordered  $\sqsupseteq$  **rel**) and a read  $e_r$  (ordered  $\sqsupseteq$  **acq**) when  $e'_w \rightarrow_{\tau}^{\text{rf}} e_r$  for  $e'_w \in$  *release-sequence*<sup>3</sup> of  $e_w$  in  $\tau$  [9,12].
- $\rightarrow_{\tau}^{\text{ithb}}$  (*Inter-thread-hb*): Inter-thread relation computed by extending  $\rightarrow_{\tau}^{\text{sw}}$  and  $\rightarrow_{\tau}^{\text{dob}}$  with  $\rightarrow_{\tau}^{\text{sb}}$ .
- $\rightarrow_{\tau}^{\text{hb}}$  (*Happens-before*): Inter-thread relation defined as  $\rightarrow_{\tau}^{\text{sb}} \cup \rightarrow_{\tau}^{\text{ithb}}$ .

The  $\rightarrow_{\tau}^{\text{hb}}$  relation along with the  $\rightarrow_{\tau}^{\text{mo}}$  and  $\rightarrow_{\tau}^{\text{rf}}$  relations (Definition 1) is used in specifying the set of six coherence conditions [12,17]:

- $\rightarrow_{\tau}^{\text{hb}}$  is irreflexive. (co-h)
- $\rightarrow_{\tau}^{\text{rf}}, \rightarrow_{\tau}^{\text{hb}}$  is irreflexive. (co-rh)
- $\rightarrow_{\tau}^{\text{mo}}, \rightarrow_{\tau}^{\text{hb}}$  is irreflexive. (co-mh)
- $\rightarrow_{\tau}^{\text{mo}}, \rightarrow_{\tau}^{\text{rf}}, \rightarrow_{\tau}^{\text{hb}}$  is irreflexive. (co-mrh)
- $\rightarrow_{\tau}^{\text{mo}}, \rightarrow_{\tau}^{\text{hb}}, \rightarrow_{\tau}^{\text{rf}^{-1}}$  is irreflexive. (co-mhi)
- $\rightarrow_{\tau}^{\text{mo}}, \rightarrow_{\tau}^{\text{rf}}, \rightarrow_{\tau}^{\text{hb}}, \rightarrow_{\tau}^{\text{rf}^{-1}}$  is irreflexive. (co-mrhi)


Additionally, *all* **sc** ordered events in a trace  $\tau$  must be related by a total order ( $\rightarrow_{\tau}^{\text{to}}$ ) that concurs with the coherence conditions. We use an irreflexive relation called *from-reads* ( $\rightarrow_{\tau}^{\text{fr}} \triangleq \rightarrow_{\tau}^{\text{rf}^{-1}}; \rightarrow_{\tau}^{\text{mo}}$ ) for ordering reads with *later* writes. Consequently,  $\rightarrow_{\tau}^{\text{to}}$  must satisfy the following condition [12,25], referred to as (**to-sc**) and formally defined in the extended version [22].

- $\forall e_1^{\text{sc}}, e_2^{\text{sc}} \in \mathcal{E}_{\tau}^{(\text{sc})}$  if  $e_1^{\text{sc}} \rightarrow_{\tau}^{\text{to}} e_2^{\text{sc}}$  then  $(e_2^{\text{sc}}, e_1^{\text{sc}}) \notin \rightarrow_{\tau}^{\text{hb}} \cup \rightarrow_{\tau}^{\text{mo}} \cup \rightarrow_{\tau}^{\text{rf}} \cup \rightarrow_{\tau}^{\text{fr}}$ ; and,
- an **sc** read (or any read with an **sc** fence  $\rightarrow_{\tau}^{\text{sb}}$  ordered before it) must not read from an **sc** write that is not *immediately*  $\rightarrow_{\tau}^{\text{to}}$  ordered before it.

Conjunction of (**coherence conditions**) and (**to-sc**) forms the sufficient condition to determine if a trace  $\tau$  is valid under *C11* (formally defined in [22]).

**HB with C11 fences.** *C11* fences form  $\rightarrow_{\tau}^{\text{ithb}}$  with other events [9,12]. A fence can be associated with the memory orders **rel**, **acq**, **ar** and **sc**. An appropriately

<sup>3</sup> *release-sequence* of  $e_w$  in  $\tau$ : maximal contiguous sub-sequence of  $\rightarrow_{\tau}^{\text{mo}}$  that starts at  $e_w$  and contains: (i) write events of  $thr(e_w)$ , (ii) rmw events of other threads [9,12].



placed fence can form  $\rightarrow_{\tau}^{\text{sw}}$  and  $\rightarrow_{\tau}^{\text{dob}}$  relation from an  $\rightarrow_{\tau}^{\text{rf}}$  relation between events of different threads (formal described in the extended version [22]).

## 5 Invalidating buggy traces with C11 fences


The key idea behind the proposed techniques is to introduce fences such that either (**coherence conditions**) or (**to-sc**) are violated. This section introduces two approaches for determining if the trace is rendered invalid with fences.

Consider  $\tau^{\text{imm}}$  of a buggy trace  $\tau$ . The candidate fences of  $\tau^{\text{imm}}$  inflate  $\rightarrow_{\tau}^{\text{sb}}$ ,  $\rightarrow_{\tau}^{\text{sw}}$ ,  $\rightarrow_{\tau}^{\text{dob}}$  and  $\rightarrow_{\tau}^{\text{ithb}}$  relations (fences do not contribute to  $\rightarrow_{\tau^{\text{imm}}}^{\text{mo}}$ ,  $\rightarrow_{\tau^{\text{imm}}}^{\text{rf}}$  and  $\rightarrow_{\tau^{\text{imm}}}^{\text{fr}}$ ). The inflated relations are denoted as  $\rightarrow_{\tau^{\text{imm}}}^{\text{sb}}$ ,  $\rightarrow_{\tau^{\text{imm}}}^{\text{sw}}$ ,  $\rightarrow_{\tau^{\text{imm}}}^{\text{dob}}$  and  $\rightarrow_{\tau^{\text{imm}}}^{\text{ithb}}$ . We propose *Weak-FenSyng* and *Strong-FenSyng* to detect the invalidity of  $\tau^{\text{imm}}$ .

**Weak-FenSyng.** Weak-FenSyng computes compositions of relations that correspond to the (**coherence conditions**). It then checks if there exist cycles in the compositions (using *Johnson's algorithm* [13]). The approach assumes the memory order **ar** for all candidate fences. Consider a buggy trace (**WRIR**) where  $x$  and  $y$  have 0 as initial values. Weak-FenSyng detects a cycle in  $\rightarrow_{\tau}^{\text{mo}}; \rightarrow_{\tau}^{\text{rf}}; \rightarrow_{\tau}^{\text{hb}}; \rightarrow_{\tau}^{\text{rf}^{-1}}$  with the addition of candidate fences  $\mathbb{F}_1^{\text{ar}}$  and  $\mathbb{F}_2^{\text{ar}}$  as shown in (**WRIR-invalidated**). This violates the condition (**co-mrhi**), thus, invalidating (**WRIR**).

**Strong-FenSyng.** This technique works with the assumption that all candidate fences have the order **sc**. Strong-FenSyng detects the infeasibility in constructing a  $\rightarrow_{\tau^{\text{imm}}}^{\text{to}}$  that adheres to (**to-sc**). In order to detect violation of (**to-sc**), Strong-FenSyng introduces a possibly reflexive relation on **sc**-ordered events of  $\tau^{\text{imm}}$ , called **sc-order** ( $\rightarrow_{\tau^{\text{imm}}}^{\text{so}}$ ). The  $\rightarrow_{\tau^{\text{imm}}}^{\text{so}}$  relation is such that a total order cannot be formed on the **sc** events of  $\tau^{\text{imm}}$  iff a cycle exists in  $\rightarrow_{\tau^{\text{imm}}}^{\text{so}}$ . All **sc** event pairs ordered by  $\rightarrow_{\tau^{\text{imm}}}^{\text{hb}}$ ,  $\rightarrow_{\tau^{\text{imm}}}^{\text{mo}}$ ,  $\rightarrow_{\tau^{\text{imm}}}^{\text{rf}}$  and  $\rightarrow_{\tau^{\text{imm}}}^{\text{fr}}$  are contained in  $\rightarrow_{\tau^{\text{imm}}}^{\text{so}}$ . Notably, pairs of **sc** events that do not have a definite order are not ordered by  $\rightarrow_{\tau^{\text{imm}}}^{\text{so}}$ . This is because if such a pair of events is involved in a cycle then we can freely flip their order and eliminate the cycle. Consider the buggy trace (**SB**),  $W^{\text{sc}}(x, 1) \rightarrow_{\tau}^{\text{to}} W^{\text{sc}}(y, 1)$  and  $W^{\text{sc}}(y, 1) \rightarrow_{\tau}^{\text{to}} W^{\text{sc}}(x, 1)$  are both valid total-orders





on the **sc** events of the trace. The set  $\rightarrow_{\tau}^{\text{so}}$  does not contain either of the two event pairs and would be empty for this example.

As a consequence, pairs of events that do not have definite total order cannot contribute to the reflexivity of  $\rightarrow_{\tau}^{\text{so}}$  and can be safely ignored. Thus,  $\rightarrow_{\tau}^{\text{so}^+} \subseteq \rightarrow_{\tau}^{\text{to}}$  for a trace  $\tau$ . Further, if a total order cannot be formed on **sc** ordered events then a corresponding cycle exists in  $\rightarrow_{\tau}^{\text{so}}$ . The observations are formally presented with supporting proofs in the extended version [22]).

Definition 3 formally presents  $\rightarrow_{\tau}^{\text{so}}$  based on the above stated considerations.

**Definition 3. sc-order** ( $\rightarrow_{\tau}^{\text{so}}$ )

$\forall e_1, e_2 \in \mathcal{E}_{\tau}$  s.t.  $(e_1, e_2) \in R$ , where  $R = \rightarrow_{\tau}^{\text{hb}} \cup \rightarrow_{\tau}^{\text{mo}} \cup \rightarrow_{\tau}^{\text{rf}} \cup \rightarrow_{\tau}^{\text{fr}}$

- if  $e_1, e_2 \in \mathcal{E}_{\tau}^{(\text{sc})}$  then  $e_1 \rightarrow_{\tau}^{\text{so}} e_2$ ; (soee)
- if  $e_1 \in \mathcal{E}_{\tau}^{(\text{sc})}$ ,  $\exists \mathbb{F}^{\text{sc}} \in \mathcal{E}_{\tau}^{\mathbb{F}(\text{sc})}$  s.t.  $e_2 \rightarrow_{\tau}^{\text{sb}} \mathbb{F}^{\text{sc}}$  then  $e_1 \rightarrow_{\tau}^{\text{so}} \mathbb{F}^{\text{sc}}$ ; (soef)
- if  $e_2 \in \mathcal{E}_{\tau}^{(\text{sc})}$ ,  $\exists \mathbb{F}^{\text{sc}} \in \mathcal{E}_{\tau}^{\mathbb{F}(\text{sc})}$  s.t.  $\mathbb{F}^{\text{sc}} \rightarrow_{\tau}^{\text{sb}} e_1$  then  $\mathbb{F}^{\text{sc}} \rightarrow_{\tau}^{\text{so}} e_2$ ; (sofe)
- if  $\exists \mathbb{F}_1^{\text{sc}}, \mathbb{F}_2^{\text{sc}} \in \mathcal{E}_{\tau}^{\mathbb{F}(\text{sc})}$  s.t.  $\mathbb{F}_1^{\text{sc}} \rightarrow_{\tau}^{\text{sb}} e_1$  and  $e_2 \rightarrow_{\tau}^{\text{sb}} \mathbb{F}_2^{\text{sc}}$  then  $\mathbb{F}_1^{\text{sc}} \rightarrow_{\tau}^{\text{so}} \mathbb{F}_2^{\text{sc}}$ . (soff)

The trace depicted in (SB) can be invalidated with strong fences as shown in (SB-inv). The **sc** events of (SB-inv) cannot be totally ordered and StrongFenSyng detects the same through a cycle in  $\rightarrow^{\text{so}}$  (formed by (soee) and (sofe)).

**Scope of FenSyng/fFenSyng.** Our work synthesizes *C11* fences and stands fundamentally different from techniques that strengthen the memory orders of events. On the one hand, **sc** fences cannot restore sequential consistency; thus, strengthening memory orders may invalidate buggy traces that the strongest *C11* fences cannot. On the other hand, strengthening may lead to sub-optimal byte-code. The difference is explained in the extended version [22].

## 6 Methodology

**Buggy traces and candidate fences.** Algorithms 1 and 2 present FenSyng and fFenSyng, respectively. The algorithms rely on an external buggy trace generator (BTG) for the buggy trace(s) of  $P$  (lines 2,8). The candidate fences are inserted (to obtain  $\tau^{\text{imm}}$ ), and the event relations are updated (lines 16-17).

**Detecting violation of trace coherence.** The algorithms detect possible violations of trace coherence conditions resulting from the candidate fences at lines 18-19 of the function `synthesisCore`. Figures (RWRW-inv-sync) and (RWRW-inv-sync-opt) represent two instances of violations of (co-rh) detected by WeakFenSyng (through a cycle  $W^{\text{rlx}}(y, 1) \rightarrow_{\tau}^{\text{rf}} R^{\text{rlx}}(y, 1) \rightarrow_{\tau}^{\text{hb}} W^{\text{rlx}}(y, 1)$ ). The candidate solutions corresponding to these cycles (which include only candidate fences) are  $\{\mathbb{F}_{12}, \mathbb{F}_{22}\}$  and  $\{\mathbb{F}_{12}\}$ . Further, for the same example, (RWRW-inv-to)



Algorithm 1: <code>FenSyng</code> ( $P$ )	Algorithm 2: <code>fFenSyng</code> ( $P$ )
<pre> 1 <math>\Phi := \top</math>; <math>\mathcal{C} := \emptyset</math> 2 forall <math>\tau \in \text{buggyTraces}(P)</math> do 3   <math>\Phi_\tau, \mathcal{C}_\tau := \text{synthesisCore}(\tau)</math> 4   <math>\Phi := \Phi \wedge \Phi_\tau</math>; <math>\mathcal{C} := \mathcal{C} \cup \mathcal{C}_\tau</math> 5 <math>\text{min}\Phi := \text{minModel}(\Phi)</math> 6 <math>\mathcal{F} := \text{assignM0}(\text{min}\Phi, \mathcal{C})</math> 7 return <math>\text{syn}(P, \mathcal{F})</math> </pre>	<pre> 8 if <math>\exists \tau \in \text{buggyTraces}(P)</math> then 9   <math>\Phi, \mathcal{C} := \text{synthesisCore}(\tau)</math> 10  <math>\text{min}\Phi := \text{minModel}(\Phi)</math> 11  <math>\mathcal{F} := \text{assignM0}(\text{min}\Phi, \mathcal{C})</math> 12  <math>P' := \text{syn}(P, \mathcal{F})</math> 13  return <code>fFenSyng</code> (<math>P'</math>) 14 else return <math>P</math> </pre>
<pre> 15 Function <code>synthesisCore</code>(<math>\tau</math>)                                     /* <math>\tau = \langle \mathcal{E}_\tau, \rightarrow_\tau^{\text{hb}}, \rightarrow_\tau^{\text{mo}}, \rightarrow_\tau^{\text{rf}} \rangle</math> */: 16   <math>\mathcal{E}_{\tau^{\text{imm}}} := \mathcal{E}_\tau \cup \text{candidateFences}(\tau)</math> 17   <math>(\rightarrow_{\tau^{\text{imm}}}^{\text{hb}}, \rightarrow_{\tau^{\text{imm}}}^{\text{mo}}, \rightarrow_{\tau^{\text{imm}}}^{\text{rf}}, \rightarrow_{\tau^{\text{imm}}}^{\text{rf}^{-1}}, \rightarrow_{\tau^{\text{imm}}}^{\text{fr}}) := \text{computeRelations}(\tau, \mathcal{E}_{\tau^{\text{imm}}})</math> 18   <math>\text{weakCycles}_\tau := \text{weakFensyng}(\tau^{\text{imm}})</math> 19   <math>\text{strongCycles}_\tau := \text{strongFensyng}(\tau^{\text{imm}})</math> 20   if <math>\text{weakCycles}_\tau = \emptyset \wedge \text{strongCycles}_\tau = \emptyset</math> then 21     return /* ABORT: cannot stop <math>\tau</math> with C11 fences */ 22   <math>\Phi_\tau := \mathcal{Q}(\text{weakCycles}_\tau \vee \text{strongCycles}_\tau)</math>; <math>\mathcal{C}_\tau := \text{weakCycles}_\tau \cup \text{strongCycles}_\tau</math> 23   return <math>\Phi_\tau, \mathcal{C}_\tau</math> </pre>	

represents a violation detected by Strong-FenSyng with the candidate solution  $\{\mathbb{F}_{12}, \mathbb{F}_{22}\}$ . The algorithms discard all candidate fences other than  $\mathbb{F}_{12}$  and  $\mathbb{F}_{22}$  from future considerations (assuming no other violations were detected). Now  $\tau$  can be invalidated as the set of cycles is nonempty (line 20).

The complexity of detecting all cycles for a trace is  $\mathcal{O}((|\mathcal{E}_\tau| + \mathbf{E}) \cdot (\mathbf{C} + 1))$  where  $\mathbf{C}$  represents the number of cycles of  $\tau$  and  $\mathbf{E}$  represents the number of pairs of events in  $\mathcal{E}_\tau$ . Note that  $\mathbf{E}$  is in  $\mathcal{O}(|\mathcal{E}_\tau|^2)$  and  $\mathbf{C}$  is in  $\mathcal{O}(|\mathcal{E}_\tau|!)$ . Thus, Weak- and Strong-Fensyng have exponential complexities in the number of traces and the number of events per trace.

**Reduction for optimality.** The algorithms use a SAT solver to determine the optimal number of candidate fences. The candidate fences from each candidate solution of  $\tau$  are conjuncted to form a SAT query. Further, to retain at least one solution corresponding to  $\tau$  the algorithms take a disjunction of the conjuncts. The SAT query is represented in the algorithm as  $\Phi_\tau := \mathcal{Q}(\text{weakCycles}_\tau \vee \text{strongCycles}_\tau)$  (line 22) and presented in Equation 1 (where  $\mathbf{W}_\tau$  and  $\mathbf{S}_\tau$  represent  $\text{weakCycles}_\tau$  and  $\text{strongCycles}_\tau$  and  $\mathbf{W}^\mathbb{F}$  and  $\mathbf{S}^\mathbb{F}$  represent the set of candidate fences in cycles  $\mathbf{W}$  and  $\mathbf{S}$  respectively). Further, `FenSyng` combines the SAT formulas corresponding to each buggy trace via conjunction (line 4), shown in Equation 2. However, note that for `fFenSyng`  $\Phi = \Phi_\tau$ .

$$\Phi_\tau = \left( \bigvee_{\mathbf{W} \in \mathbf{W}_\tau} \bigwedge_{\mathbb{F}_w \in \mathbf{W}^\mathbb{F}} \mathbb{F}_w \right) \vee \left( \bigvee_{\mathbf{S} \in \mathbf{S}_\tau} \bigwedge_{\mathbb{F}_s \in \mathbf{S}^\mathbb{F}} \mathbb{F}_s \right) \quad (1) \qquad \Phi = \bigwedge_{\tau \in \text{BT}} \Phi_\tau \quad (2)$$

We use a SAT solver to compute the *min-model* ( $\text{min}\Phi$ ) of the query  $\Phi$  (lines 5,10). For instance, the query for (`RWRW-bt`) is  $\Phi = (\mathbb{F}_{12}) \vee (\mathbb{F}_{12} \wedge \mathbb{F}_{22}) \vee (\mathbb{F}_{12} \wedge \mathbb{F}_{22})$  and min-model,  $\text{min}\Phi = \{\mathbb{F}_{12}\}$ . The solution to the SAT query returns the smallest set of fences to be synthesized.

The complexity of constructing the query  $\Phi_\tau$  is  $\mathcal{O}(\mathbf{C}\cdot\mathbf{F})$ , where  $\mathbf{C}$  is the number of cycles per trace and  $\mathbf{F}$  is the number of fences per cycle. The structure of the query  $\Phi$  corresponds to the *Head-cycle-free* (HCF) class of CNF theories; hence, the min-model computation falls in the FP complexity class [8].

**Determining optimal memory orders of fences.** The set  $\min\Phi$  gives a sound solution that is optimal only in the number of fences. The function `assignMO` (lines 6,11) assigns the weakest memory order to the fences in  $\min\Phi$  that is sound. Let `min-cycles` represent a set of cycles such that every candidate fence in the cycles belongs to  $\min\Phi$ . The `assignMO` function computes memory order for fences of `min-cycles` of each trace as follows: If a cycle  $c \in \text{min-cycles}$  is detected, then its fences must form a  $\rightarrow_{\tau^{\text{imm}}}^{\text{sw}}$  or  $\rightarrow_{\tau^{\text{imm}}}^{\text{dob}}$  with an event of  $\tau^{\text{imm}}$  (since, candidate fences only modify  $\rightarrow_{\tau}^{\text{sb}}$ ,  $\rightarrow_{\tau}^{\text{sw}}$  and  $\rightarrow_{\tau}^{\text{dob}}$ ). Let  $R = \rightarrow_{\tau^{\text{imm}}}^{\text{sw}} \cup \rightarrow_{\tau^{\text{imm}}}^{\text{dob}}$ . The scheme to compute fence types is as follows:

- If a fence  $\mathbb{F}$  in a weak cycle  $c$  is related to an event  $e$  of  $c$  by  $R$  as  $eR\mathbb{F}$ , then  $\mathbb{F}$  is assigned the memory order `acq`;
- if an event  $e$  in  $c$  is related to  $\mathbb{F}$  as  $\mathbb{F}Re$  then  $\mathbb{F}$  is assigned `rel`;
- if events  $e, e'$  of  $c$  are related to  $\mathbb{F}$  as  $eR\mathbb{F}Re'$  then  $\mathbb{F}$  is assigned `ar`.
- All the fences in a strong cycle are assigned the memory order `sc`.

Consider a cycle  $c: e \rightarrow_{\tau^{\text{imm}}}^{\text{sb}} \mathbb{F}_1 \rightarrow_{\tau^{\text{imm}}}^{\text{sw}} \mathbb{F}_2 \rightarrow_{\tau^{\text{imm}}}^{\text{sw}} \mathbb{F}_3 \rightarrow_{\tau^{\text{imm}}}^{\text{sb}} e' \rightarrow_{\tau^{\text{imm}}}^{\text{rf}} e$  representing a violation of  $\rightarrow_{\tau^{\text{imm}}}^{\text{rf}}$ ;  $\rightarrow_{\tau^{\text{imm}}}^{\text{hb}}$  irreflexivity (condition `(co-rh)`). According to the scheme discussed above, the fences  $\mathbb{F}_1, \mathbb{F}_2$  and  $\mathbb{F}_3$  are assigned the memory orders `rel`, `ar` and `acq` respectively and  $wt(c) = 4$  (defined in §3).

Further, `assignMO` iterates over all buggy traces and detects the sound weakest memory order for each fence across all traces as follows. Assume a cycle  $c_1$  in  $\tau_1^{\text{imm}}$  and a cycle  $c_2$  in  $\tau_2^{\text{imm}}$ . The function computes a union of the fences of  $\tau_1$  and  $\tau_2$  while choosing the stronger memory order for each fence that is present in both the cycles. In doing so, both  $\tau_1$  and  $\tau_2$  are invalidated. Further, when two candidate solutions have the same set of fences, the function selects the one with the lower weight.

Consider the cycles of buggy traces  $\tau_1$  and  $\tau_2$  shown in `(candidate-fences)`. Let  $\min\Phi = \{\mathbb{F}_1, \mathbb{F}_2, \mathbb{F}_3\}$ . The memory orders of the fences for each trace are shown with superscripts and the weights of the cycle  $\tau_1c_1$ ,  $\tau_1c_2$  and  $\tau_2c_1$  are written against the name of the cycles. The candidate solutions  $\tau_1c_1$  and  $\tau_1c_2$  are combined with  $\tau_2c_1$  to form  $\tau_{12}c_{11}$  and  $\tau_{12}c_{21}$  of weights 5 and 4, respectively. The solution  $\tau_{12}c_{11}$  is of higher weight and is discarded. In  $\tau_{12}c_{21}$ , the optimal memory orders `rel`, `acq` and `ar` are assigned to fences  $\mathbb{F}_1, \mathbb{F}_2$  and  $\mathbb{F}_3$ , respectively. It is possible that `min-cycles` may contain fences originally in  $P$ . If the process discussed above computes a stronger memory order for a program fence than its original order in  $P$ , then the technique strengthens the memory order of the fence to the computed order. Note that this reasoning across traces does not occur in `fFenSyng` as it considers only one trace at a time.

Determining the optimal memory orders has a complexity in  $\mathcal{O}(\text{BT}\cdot\mathbf{C}\cdot\mathbf{F} + \mathbf{M}^{\text{BT}})$ , where BT is the number of buggy traces of  $P$ ,  $\mathbf{C}$  and  $\mathbf{F}$  are defined as before, and  $\mathbf{M}$  is the number of min-cycles per trace.

$\tau_1 c_1(4): \mathbb{F}_1^{\text{ar}} \wedge \mathbb{F}_2^{\text{ar}}$	cycles of $\tau_1$	cycles in $\tau_1$ ( $\mathbf{C}_{\tau_1}$ ): $\{\mathbb{F}_1, \mathbb{F}_2, e_1\}$ and $\{\mathbb{F}_1, \mathbb{F}_3, \mathbb{F}_4\}$ $\Phi_{\tau_1} = (\mathbb{F}_1 \wedge \mathbb{F}_2) \vee (\mathbb{F}_1 \wedge \mathbb{F}_3 \wedge \mathbb{F}_4)$
$\tau_1 c_2(4): \mathbb{F}_1^{\text{rel}} \wedge \mathbb{F}_2^{\text{acq}} \wedge \mathbb{F}_3^{\text{ar}}$		
$\tau_2 c_1(3): \mathbb{F}_1^{\text{rel}} \wedge \mathbb{F}_2^{\text{acq}} \wedge \mathbb{F}_3^{\text{acq}}$		
$\tau_{12} c_{11}(5): \mathbb{F}_1^{\text{ar}} \wedge \mathbb{F}_2^{\text{ar}} \wedge \mathbb{F}_3^{\text{acq}}$	candidate-fences	cycles in $\tau_2$ ( $\mathbf{C}_{\tau_2}$ ): $\{\mathbb{F}_3, \mathbb{F}_4\}$ $\Phi_{\tau_2} = (\mathbb{F}_3 \wedge \mathbb{F}_4)$ 3-fence
$\tau_{12} c_{21}(4): \mathbb{F}_1^{\text{rel}} \wedge \mathbb{F}_2^{\text{acq}} \wedge \mathbb{F}_3^{\text{ar}}$		

In our experimental observation (refer to §7), the number of buggy traces analyzed by **fFenSyng** is significantly less than  $|\text{BT}|$ . Therefore, in practice, the complexity of various steps of **fFenSyng** that are dependent on BT reduces exponentially by a factor of  $|\text{BT}|$ .

**Nonoptimality of fFenSyng.** Consider the example (3-fence). It shows cycles in two buggy traces  $\tau_1$  and  $\tau_2$  of an input program. **FenSyng** provides the formula  $\Phi_{\tau_1} \wedge \Phi_{\tau_2}$  to the SAT solver and the optimal solution obtained is  $(\mathbb{F}_1 \wedge \mathbb{F}_3 \wedge \mathbb{F}_4)$ . However, **fFenSyng** considers the formula  $\Phi_{\tau_1}$  and  $\Phi_{\tau_2}$  in separate iterations and may return a nonoptimal result  $(\mathbb{F}_1 \wedge \mathbb{F}_2) \wedge (\mathbb{F}_3 \wedge \mathbb{F}_4)$ .

We prove the soundness of **fFenSyng** and **FenSyng** with Theorems 1 and 2 respectively and the optimality of **FenSyng** with Theorem 3. The theorems are formally presented with proofs in the extended version [22].

**Theorem 1.** *fFenSyng is sound. If a buggy trace  $\tau$  of  $P$  can be invalidated using C11 fences then fFenSyng will invalidate  $\tau$ .*

**Theorem 2.** *FenSyng is sound. If a buggy program  $P$  can be fixed using C11 fences, then FenSyng will invalidate all buggy traces of  $P$ .*

**Theorem 3.** *FenSyng is optimal. FenSyng synthesizes optimal number of fences with optimal memory orders.*

## 7 Implementation and Results

**Implementation details.** The techniques are implemented in Python. Weak-FenSyng and Strong-FenSyng use *Johnson's* cycle detection algorithm in the *networkx* library. We use Z3 theorem prover to find the *min-model* of SAT queries. As a BTG, we use CDSChecker [20], an open-source model checker, for the following reasons;

1. CDSChecker supports the C11 semantics. Most other techniques are designed for a variant [15] or subset [1,3,23] of C11.
2. CDSChecker returns buggy traces along with the corresponding  $\rightarrow_{\tau}^{\text{hb}}$ ,  $\rightarrow_{\tau}^{\text{rf}}$  and  $\rightarrow_{\tau}^{\text{mo}}$  relations.
3. CDSChecker does not halt at the detection of the first buggy trace; instead, it continues to provide all buggy traces as required by **FenSyng**.

To bridge the gap between CDSChecker's output and our requirements, we modify CDSChecker's code to accept program location as an attribute of the program events and to halt at the first buggy trace when specified. **FenSyng** and

**Table 1.** Litmus Testing Summary

*Litmus Tests Summary*

Tests	min-BT	max-BT	avg-BT	min-syn	max-syn	avg-syn	min-str	max-str	avg-str
1389	1	9	1.05	1	4	2.25	0	0	0

BT: #buggy traces, syn: #fences synthesized, str: # fences strengthened  
min: minimum, max: maximum, avg: average

*Results Summary*

	completed (syn+no fix)	TO	NO	Tbtg (total)	TF (total)	Ttotal
<b>FenSyng</b>	1333 (1185+148)	56	0	50453.19	36896.06	87266.09
<b>fFenSyng</b>	1355 (1207+148)	34	0	30703.71	49068.61	79772.32

Times in seconds.

TO: 15min for BTG + 15min for technique

Tbtg: Time of BTG, TF: Time of **FenSyng** or **fFenSyng**, Ttotal: Tbtg+TF

**fFenSyng** are available as an open-source tool that performs fence synthesis for *C11* programs at: <https://github.com/singhsanjana/fensyng>.

**Experimental setup.** The experiments were performed on an Intel(R) Xeon(R) CPU E5-1650 v4 @ 3.60GHz with 32GB RAM and 32 cores. We collected a set of 1389 litmus tests of buggy *C11* input programs (borrowed from Tracer [3]) to validate the correctness of **FenSyng** and **fFenSyng** experimentally. We study the performance of **FenSyng** and **fFenSyng** on a set of benchmarks borrowed from previous works on model checking under *C11* and its variants [1,3,20,23].

**Experimental validation.** The summary of the 1389 litmus tests is shown under *Litmus Tests Summary*, Table 1. The number of buggy traces for the litmus tests ranged between 1-9 with an average of 1.05, while the number of fences synthesized ranged between 2-4. None of the litmus tests contained fences in the input program. Hence, no fences were strengthened in any of the tests.

We present the results of **FenSyng** and **fFenSyng** under *Result Summary*, Table 1. The results have been averaged over five runs for each test. **fFenSyng** timed out (column ‘TO’) on a fewer number of tests (34 tests) in comparison to **FenSyng** (56 tests). The techniques could not fix 148 tests with *C11* fences (‘no fix’). The column ‘NO’ represents the number of tests where the fences synthesized or strengthened is nonoptimal. To report the values of ‘NO’, we conducted a sanity test on the fixed program as follows: we create versions  $P_1, \dots, P_k$  of the fixed program  $P^{fx}$  s.t. in each version, one of the fences of  $P^{fx}$  is either weakened or eliminated. Each version is then tested separately on BTG. The sanity check is successful if a buggy trace is returned for each version.

**Performance analysis.** We contrast the performance of the techniques using a set of benchmarks that produce buggy traces under *C11*. The results are averaged over five runs. Table 2 reports the results where ‘#BT’ shows the number of buggy traces, ‘iter’ shows the minimum:maximum number of iterations performed by **fFenSyng** over the five runs and, ‘FTo’ and ‘BTo’ represent **FenSyng/fFenSyng** time-out and BTG time-out, respectively (set to 15 minutes each). A ‘?’ in ‘#BT’ signifies that BTG could not scale for the test, so the number of buggy traces is unknown. The column (‘syn+str’) under **fFenSyng** reports the minimum:maximum number of fences synthesized and/or strengthened. We

**Table 2.** Comparative performance analysis


Id	Name	#BT	FenSyng				fFenSyng				
			syn+str	Tbtg	TF	Ttotal	iter	syn+str	Tbtg	TF	Ttotal
1	peterson(2,2)	30	1+0	2.63	54.31	56.94	1:1	1:1+0:0	0.18	2.07	2.25
2	peterson(2,3)	198	1+0	29.96	594.34	624.3	1:1	1:1+0:0	0.53	3.58	4.11
3	peterson(4,5)	?	-	-	FTo	-	1:1	1:1+0:0	397.51	21.07	418.58
4	peterson(5,5)	?	-	BTo	-	-	1:1	1:1+0:0	BTo	31.52	*931.52
5	barrier(5)	136	1+0	1.09	207.74	208.83	1:1	1:1+0:0	0.13	1.40	1.53
6	barrier(10)	416	1+0	3.37	565.44	568.81	1:1	1:1+0:0	0.2	2.70	2.9
7	barrier(100)	31106	-	-	FTo	-	1:1	1:1+0:0	34.2	198.54	232.74
8	barrier(150)	?	-	-	FTo	-	1:1	1:1+0:0	117.09	399.20	516.29
9	barrier(200)	-	-	-	-	-	-	-	-	FTo	-
10	store-buffer(2)	6	2+0	0.08	0.91	0.99	1:1	2:2+0:0	0.04	0.05	0.09
11	store-buffer(4)	20	2+0	1.61	195.35	196.96	1:1	2:2+0:0	1.20	0.05	1.25
12	store-buffer(5)	30	-	-	FTo	-	1:1	2:2+0:0	14.07	0.22	14.29
13	store-buffer(6)	42	-	-	FTo	-	1:1	2:2+0:0	171.09	0.15	171.24
14	store-buffer(10)	?	-	BTo	-	-	1:1	2:2+0:0	BTo	0.05	*900.05
15	dekker(2)	54	2+0	0.17	0.27	0.44	1:1	2:2+0:0	0.26	0.04	0.3
16	dekker(3)	1596	-	-	FTo	-	1:1	2:2+0:0	586.46	1.34	587.8
17	dekker-fen(2,3)	54	1+1	0.15	0.29	0.44	1:1	1:1+1:1	0.25	0.05	0.3
18	dekker-fen(3,2)	730	-	-	FTo	-	1:1	1:1+1:1	159.84	5.56	165.4
19	dekker-fen(3,4)	3076	-	BTo	-	-	1:1	1:1+1:1	BTo	6.06	*906.06
20	burns(1)	36	-	-	FTo	-	7:8	8:10+2:2	0.61	4.69	5.3
21	burns(2)	10150	-	-	FTo	-	6:7	8:10+0:1	71.53	554.6	626.13
22	burns(3)	?	-	BTo	-	-	-	-	-	FTo	-
23	burns-fen(2)	100708	-	-	FTo	-	5:7	4:6+3:3	329.41	43.96	373.37
24	burns-fen(3)	?	-	BTo	-	-	5:7	4:6+3:3	BTo	70.14	*970.14
25	linuxrwlocks(2,1)	10	-	-	FTo	-	1:1	2:2+0:0	0.13	0.12	0.25
26	linuxrwlocks(3,8)	353	-	-	FTo	-	2:2	3:4+0:0	686.52	0.41	*686.93
27	seqlock(2,1,2)	500	-	-	FTo	-	1:1	1:1+0:0	341.54	2.38	343.92
28	seqlock(1,2,2)	592	-	-	FTo	-	1:2	1:2+0:0	119.88	27.69	147.57
29	seqlock(2,2,3)	?	-	BTo	-	-	1:2	1:2+0:0	BTo	88.52	988.52*
30	bakery(2,1)	6	1+0	0.25	25.42	2.88	1:1	1:1+0:0	0.07	0.18	0.25
31	bakery(4,3)	7272	-	-	FTo	-	1:1	1:1+0:0	166.11	5.68	171.79
32	bakery(4,4)	50402	-	-	FTo	-	1:1	1:1+0:0	BTo	18.17	918.17*
33	lamport(1,1,2)	1	No fix.	0.06	0.05	0.11	1:1	No fix.	0.04	0.05	0.09
34	lamport(2,2,1)	1	No fix.	411.94	0.05	411.99	1:1	No fix.	53.34	0.05	53.39
35	lamport(2,2,3)	?	-	BTo	-	-	1:1	No fix.	389.77	0.05	389.82
36	flipper(5)	297	2+0	6.22	254.18	260.40	1:1	2+0	2.51	0.02	2.53
37	flipper(7)	4493	-	-	FTo	-	1:1	2+0	119.21	0.02	119.23
38	flipper(10)	?	-	-	FTo	-	1:1	2+0	BTo	0.03	900.03*

Tbtg: Time of BTG, TF: Time of technique (FenSyng or fFenSyng), Ttotal: Tbtg+TF

add a ‘\*’ against the time when BTG timed out in detecting that the fixed program has no more buggy traces.

The performance of FenSyng and fFenSyng is diagrammatically contrasted in Figure 1. It is notable that fFenSyng significantly outperforms FenSyng in terms of the time of execution and scalability and adds extra fences in only 7 tests with an average of 1.57 additional fences. With the increase in the number of buggy traces, an exponential rise in FenSyng’s time leading to FTo was observed; except in cases 12, 13, 20, and 25, where FenSyng times out with as low as 10 traces. The tests time-out in Johnson’s cycle detection due to a high density of the number of related events or the number of cycles.

fFenSyng analyzes a remarkably smaller number of buggy traces (‘iter’) in comparison with ‘#BT’ ( $\leq 2$  traces for  $\sim 85\%$  of tests). We conclude that a solution corresponding to a single buggy trace fixes more than one buggy traces. As a result, fFenSyng can scale to tests with thousands of buggy traces and



\* represents the remaining Test IDs (tests that timeout for both FenSyng and fFenSyng)

**Fig. 1.** Performance comparison between FenSyng and fFenSyng

we witness an average speedup of over 67x, with over 100x speedup in  $\sim 41\%$  of tests, against FenSyng.

*Interesting cases.* Consider test 16, where BTG times out in 3/5 runs and completes in  $\sim 100$ s in the remaining 2 runs. A fence is synthesized between two events,  $e_1$  and  $e_2$ , that are inside a loop. Additionally,  $e_1$  is within a condition. Depending on where the fence is synthesized (within the condition or outside it), BTG either runs out of time or finishes quickly. Similarly, BTG for test 26 times out in 3/5 runs. However, the reason here is the additional nonoptimal fences synthesized that increase the analysis overhead of the chosen BTG (CDSChecker).

Note that, for most benchmarks, fFenSyng’s scalability is limited by BTo and observably fFenSyng’s time is much lesser than FTo for such cases. Therefore, an alternative BTG would significantly improve fFenSyng’s performance.

## 8 Related Work

The literature on fence synthesis is rich with techniques targeting the x86-TSO [2,5,6,7,10] and sparc-PSO [4,18] memory models or both [14,16,19]. The work in [10] and [16] perform fence synthesis for ARMv7 and RMO memory models. The works in [4,7,11] are proposed for Power memory model, where [11] also supports IA-32 memory model.

Most fence synthesis techniques introduce additional ordering in the program events with the help of fences [5,6,7,10,11,14,18,19,24]. However, the axiomatic definition of ordering varies with memory models. As a consequence, most existing techniques (such as those for TSO and PSO) may not detect C11 buggy traces due to a strong implicit ordering. While the techniques [6,7,24] are parametric in or oblivious to the memory model, they introduce ordering between *pairs* of events that is *globally visible* (to all threads). Such an ordering constraint is restrictive for weaker models such as C11 that may require ordering on a *set* of events that may be *conditionally visible* to a thread. Similarly, [14] proposes a bounded technique applicable to any memory model that supports interleaving with reordering. Program outcomes under C11 may not be feasible under such a model. Moreover, any existing technique, cannot fix a C11 input program while conserving its portability.

Some earlier works such as [2,7,11] synthesize fences to restrict outcomes to SC or its variant for store-buffering [5]. Most fence synthesis techniques

[4,14,16,18,19] attempt to remove traces violating a safety property specification under their respective axiomatic definition of memory model. Various works [4,5,10,14,16,19,24] perform optimal fence synthesis where the optimality (in the absence of types of fences) is simply defined as the smallest set of fences. Technique [6] assigns weights to various types of fences (similar to our work) and defines optimality on the summation of fence weights of candidate solutions. However, their definition of optimality is incomparable with ours, and no prior work establishes the advantage of one definition over the other.

Lastly, a recent technique [21] fixes a buggy *C11* program by strengthening memory access events instead of synthesizing fences.

## 9 Conclusion and Future Work

This paper proposed the first fence synthesis techniques for *C11* programs: an optimal (**FenSyng**) and a near-optimal (**fFenSyng**). The work also presented theoretical arguments that showed the correctness of the synthesis techniques. The experimental validation demonstrated the effectiveness of **fFenSyng** vis-à-vis optimal **FenSyng**. As part of future work, we will investigate extending the presented methods (i) to support richer constructs such as locks and (ii) to include strengthening memory accesses to fix buggy traces.

## References

1. Abdulla, P.A., Arora, J., Atig, M.F., Krishna, S.: Verification of programs under the release-acquire semantics. In: Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation. pp. 1117–1132 (2019)
2. Abdulla, P.A., Atig, M.F., Chen, Y.F., Leonardsson, C., Rezine, A.: Automatic fence insertion in integer programs via predicate abstraction. In: International Static Analysis Symposium. pp. 164–180. Springer (2012)
3. Abdulla, P.A., Atig, M.F., Jonsson, B., Ngo, T.P.: Optimal stateless model checking under the release-acquire semantics. Proceedings of the ACM on Programming Languages 2(OOPSLA), 1–29 (2018)
4. Abdulla, P.A., Atig, M.F., Lång, M., Ngo, T.P.: Precise and sound automatic fence insertion procedure under pso. In: International Conference on Networked Systems. pp. 32–47. Springer (2015)
5. Abdulla, P.A., Atig, M.F., Ngo, T.P.: The best of both worlds: Trading efficiency and optimality in fence insertion for tso. In: European Symposium on Programming Languages and Systems. pp. 308–332. Springer (2015)
6. Alglave, J., Kroening, D., Nival, V., Poetzl, D.: Don’t sit on the fence. In: International Conference on Computer Aided Verification. pp. 508–524. Springer (2014)
7. Alglave, J., Maranget, L., Sarkar, S., Sewell, P.: Fences in weak memory models. In: International Conference on Computer Aided Verification. pp. 258–272. Springer (2010)
8. Angiulli, F., Ben-Eliyahu-Zohary, R., Fassetti, F., Palopoli, L.: On the tractability of minimal model computation for some cnf theories. Artificial Intelligence 210, 56–77 (2014)



9. Batty, M., Owens, S., Sarkar, S., Sewell, P., Weber, T.: Mathematizing c++ concurrency. *ACM SIGPLAN Notices* 46(1), 55–66 (2011)
10. Bender, J., Lesani, M., Palsberg, J.: Declarative fence insertion. *ACM SIGPLAN Notices* 50(10), 367–385 (2015)
11. Fang, X., Lee, J., Midkiff, S.P.: Automatic fence insertion for shared memory multiprocessing. In: *Proceedings of the 17th annual international conference on Supercomputing*. pp. 285–294 (2003)
12. ISO/IEC-JTC1/SC22/WG21: Programming languages — C++ (2013), <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2013/n3690.pdf>
13. Johnson, D.B.: Finding all the elementary circuits of a directed graph. *SIAM Journal on Computing* 4(1), 77–84 (1975)
14. Joshi, S., Kroening, D.: Property-driven fence insertion using reorder bounded model checking. In: *International Symposium on Formal Methods*. pp. 291–307. Springer (2015)
15. Kokologiannakis, M., Raad, A., Vafeiadis, V.: Model checking for weakly consistent libraries. In: *Proceedings of the 40th ACM SIGPLAN Conference on Programming Language Design and Implementation*. pp. 96–110 (2019)
16. Kuperstein, M., Vechev, M., Yahav, E.: Automatic inference of memory fences. *ACM SIGACT News* 43(2), 108–123 (2012)
17. Lahav, O., Vafeiadis, V., Kang, J., Hur, C.K., Dreyer, D.: Repairing sequential consistency in c/c++ 11. vol. 52, pp. 618–632. ACM New York, NY, USA (2017)
18. Linden, A., Wolper, P.: A verification-based approach to memory fence insertion in pso memory systems. In: *International Conference on Tools and Algorithms for the Construction and Analysis of Systems*. pp. 339–353. Springer (2013)
19. Meshman, Y., Dan, A., Vechev, M., Yahav, E.: Synthesis of memory fences via refinement propagation. In: *International Static Analysis Symposium*. pp. 237–252. Springer (2014)
20. Norris, B., Demsky, B.: Cdschecker: checking concurrent data structures written with c/c++ atomics. In: *Proceedings of the 2013 ACM SIGPLAN international conference on Object oriented programming systems languages & applications*. pp. 131–150 (2013)
21. Oberhauser, J., Chehab, R.L.d.L., Behrens, D., Fu, M., Paolillo, A., Oberhauser, L., Bhat, K., Wen, Y., Chen, H., Kim, J., et al.: Vsync: push-button verification and optimization for synchronization primitives on weak memory models. In: *Proceedings of the 26th ACM International Conference on Architectural Support for Programming Languages and Operating Systems*. pp. 530–545 (2021)
22. Singh, S., Sharma, D., Jaju, I., Sharma, S.: Fence synthesis under the c11 memory model. *arXiv preprint arXiv:2208.00285* (2022)
23. Singh, S., Sharma, D., Sharma, S.: Dynamic verification of c11 concurrency over multi copy atomics. In: *2021 International Symposium on Theoretical Aspects of Software Engineering (TASE)*. pp. 39–46. IEEE (2021)
24. Taheri, M., Pourdamghani, A., Lesani, M.: Polynomial-time fence insertion for structured programs. In: *33rd International Symposium on Distributed Computing (DISC 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik (2019)
25. Vafeiadis, V., Balabonski, T., Chakraborty, S., Morisset, R., Zappa Nardelli, F.: Common compiler optimisations are invalid in the c11 memory model and what we can do about it. In: *Proceedings of the 42Nd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*. pp. 209–220 (2015)