

Control Flow, Scope, and Advanced Lifetime Cases

Ivan Murashko

Contents

1	Memory operations	2
1.1	Store operation dispatch	2
1.2	Coroutine task stores	5
1.3	Lambda capture stores	6
1.4	Load operation and dereference	6
2	Control flow analysis	8
2.1	Loop analysis	10
2.2	Switch statement analysis	13
2.3	The JOIN operation	16
3	Scope and lifetime management	18
3.1	Lexical scope tracking	18
3.2	The KILL operation	18
4	Coroutines and async code	22
4.1	Await operation handling	22
4.2	Coroutine task lifetime tracking	23
4.3	Temporary vs persistent tasks	25
4.4	Lambda captures by reference	26
5	Return value safety	27
6	Implementation patterns and best practices	28
6.1	Semantic helper methods	29
6.2	Composable building blocks	29
6.3	API symmetry	29
6.4	Early-exit pattern	30

7	Limitations and future work	31
8	Lessons learned	32

This article finishes the ClangIR lifetime analysis series. We will use the points-to model from [Building the LifetimeCheck Pass](#) and the moved-from state tracking from [AST Semantics and Use-After-Move Detection](#). The remaining question is how the same analysis behaves when the program has stores, loads, branches, loops, scopes, coroutines, and return values.

1 Memory operations

Store and load operations are the workhorses of lifetime tracking. Nearly every assignment, initialization, and value use flows through `checkStore()` and `checkLoad()`. Understanding how these operations dispatch to specialized handlers is crucial for understanding the complete picture of lifetime analysis.

1.1 Store operation dispatch

When a `StoreOp` is encountered (representing any assignment or initialization in C++), the pass must determine what kind of store it is and handle it appropriately. Different store patterns have different lifetime semantics.

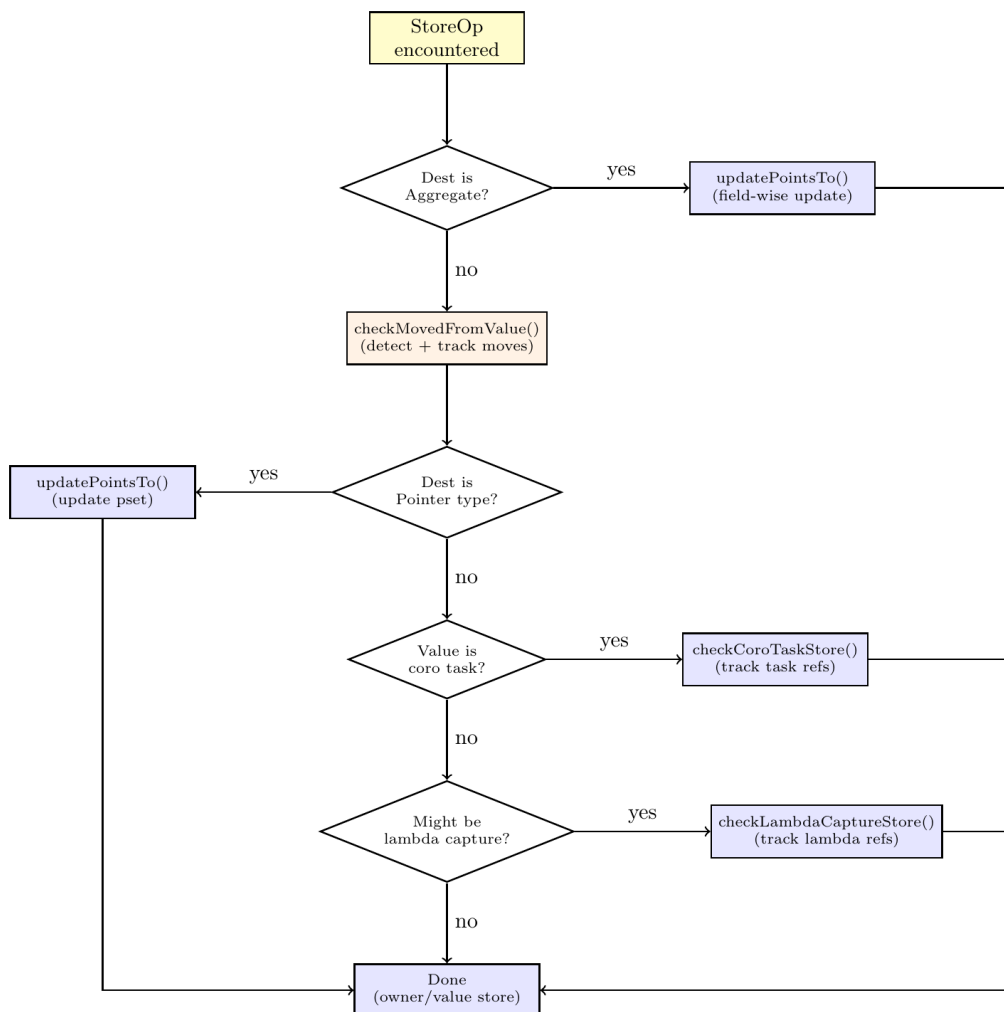


Figure 1: Store operation dispatch decision flow in `checkStore()`

The store dispatch logic follows this order:

1. **Aggregate stores** — If the destination is an aggregate type (struct/class tracked in the `aggregates` map), call `updatePointsTo()` to update individual field psets. This handles bulk initialization like `MyStruct s = {ptr1, ptr2};`.
2. **Moved-from value check** — *Always* call `checkMovedFromValue()` to detect if the source value is moved-from (error) and to track rvalue initialization (state update). This is the dual-purpose check explained in [the previous article](#).

3. **Pointer stores** — If the destination is a pointer type (in the `ptrs` set), call `updatePointsTo()` to update its pset. This is the standard path for assignments like `ptr = &x;`.
4. **Coroutine task stores** — If the stored value is a coroutine task temporary (tracked in `currScope->localTempTasks`), call `checkCoroutineTaskStore()` to bind local values used in task initialization to the task's pset.
5. **Lambda capture stores** — Otherwise, call `checkLambdaCaptureStore()` to handle potential lambda captures by reference, which can create dangling references if the captured local goes out of scope.
6. **Owner/Value stores** — If none of the above apply, the store is to an owner or value type and requires no special handling (the value is copied/moved as normal).

Implementation pattern from `LifetimeCheck.cpp`, simplified:

```

1 void LifetimeCheckPass::checkStore(StoreOp storeOp) {
2     auto addr = storeOp.getAddr();
3
4     // Handle aggregate stores (field-wise update)
5     if (aggregates.count(addr)) {
6         auto data = storeOp.getValue();
7         if (data.getDefiningOp<cir::ConstantOp>()) {
8             updatePointsTo(addr, data, data.getLoc());
9         }
10    }
11    return;
12
13    // Always check for moved-from values (dual purpose: detect +
14    ↪ track)
15    checkMovedFromValue(storeOp);
16
17    // Special handling for non-pointer types
18    if (!ptrs.count(addr)) {
19        if (currScope->localTempTasks.count(storeOp.getValue()))
20            checkCoroutineTaskStore(storeOp);
21        else
22            checkLambdaCaptureStore(storeOp);
23    }
24    return;
25 }

```

```

25 // Standard pointer store: update points-to set
26 updatePointsTo(addr, storeOp.getValue(),
    ↪ storeOp.getValue().getLoc());
27 }

```

Key insight: The *order* of checks matters. Aggregate stores are decomposed first and may return early. For the remaining store cases, the pass checks for a moved-from value before it handles pointer updates, coroutine task stores, or lambda captures, so use-after-move is detected before the store is interpreted as a normal assignment.

1.2 Coroutine task stores

Coroutine tasks are special because they can capture references to local variables that may outlive the coroutine frame. When a task is initialized, the pass tracks which local values are passed as arguments, and binds them to the task's pset.

```

1 void LifetimeCheckPass::checkCoroTaskStore(StoreOp storeOp) {
2     auto taskTmp = storeOp.getValue();
3     auto taskAddr = storeOp.getAddr();
4
5     // Pattern: %tmp_task = cir.call @coroutine_call(%arg0, %arg1,
    ↪ ...)
6     //         cir.store %tmp_task, %task
7     //
8     // Bind local values used as arguments to pset(task)
9     if (auto call = taskTmp.getDefiningOp<cir::CallOp>()) {
10        bool potentialTaintedTask = false;
11        for (auto arg : call.getArgOperands()) {
12            auto alloca = arg.getDefiningOp<cir::AllocaOp>();
13            if (alloca && currScope->localValues.count(alloca)) {
14                // Task now depends on this local value
15                getPmap()[taskAddr].insert(State::getLocalValue(alloca));
16                potentialTaintedTask = true;
17            }
18        }
19    }
20 }

```

This ensures that when a local variable goes out of scope (triggering KILL), any tasks that reference it are also invalidated. Example:

```

1 Task<int> createTask() {
2     int local = 42;
3     auto task = asyncCompute(&local); // task captures local by ref
4     // End of scope: KILL(local) also invalidates task
5     return task; // WARNING: task references destroyed local
6 }

```

1.3 Lambda capture stores

Lambda captures by reference can similarly create dangling references. When a local value is stored into a lambda capture field, the pass binds that local to the lambda's pset:

```

1 void LifetimeCheckPass::checkLambdaCaptureStore(StoreOp storeOp) {
2     auto localByRefAddr = storeOp.getValue();
3     auto lambdaCaptureAddr = storeOp.getAddr();
4
5     if (!localByRefAddr.getDefiningOp<cir::AllocaOp>())
6         return;
7     auto lambdaAddr = getLambdaFromMemberAccess(lambdaCaptureAddr);
8     if (!lambdaAddr)
9         return;
10
11     // Bind captured local to lambda's pset
12     if (currScope->localValues.count(localByRefAddr)
13         getPmap()[lambdaAddr].insert(State::getLocalValue(localByRefAddr)
14         ↪ r));
14 }

```

Example:

```

1 auto makeLambda() {
2     int local = 42;
3     auto lambda = [&local]() { return local + 1; };
4     // lambda captures local by reference
5     // End of scope: KILL(local) invalidates lambda
6     return lambda; // WARNING: lambda references destroyed local
7 }

```

1.4 Load operation and dereference

Load operations are simpler than stores—they primarily check validity when dereferencing pointers. The key insight is the `isDeref` flag, which distin-

guishes between loading a pointer value (safe) and dereferencing it (requires validity check).

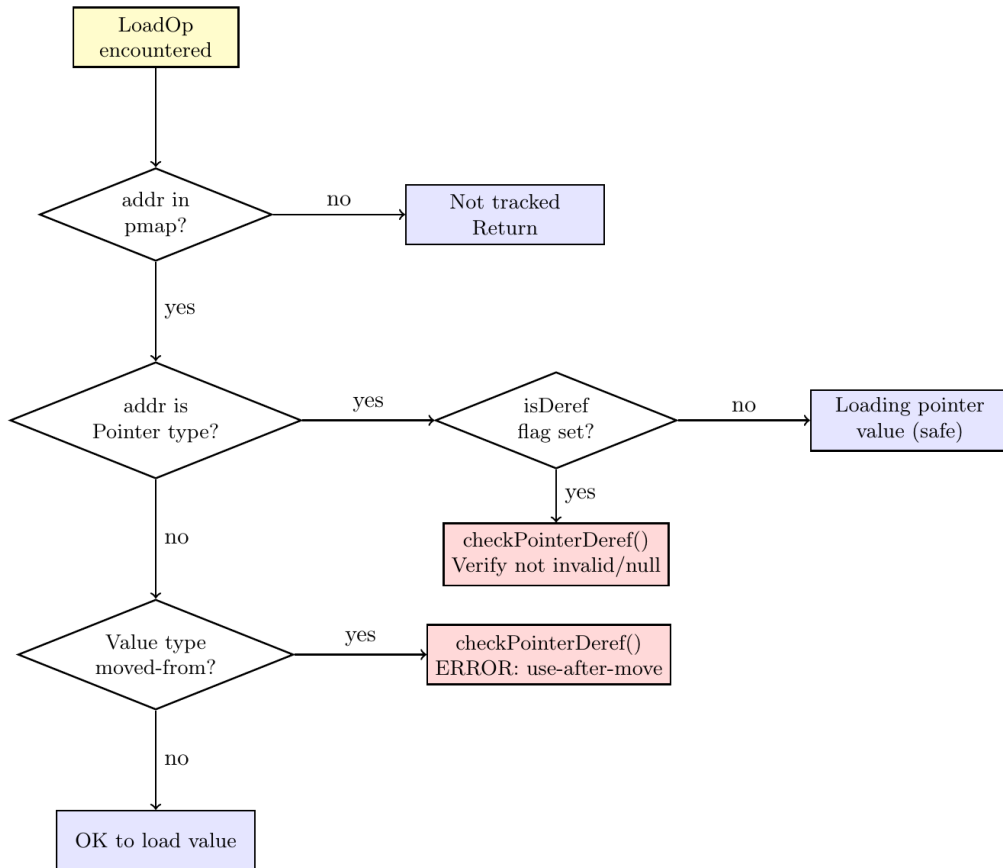


Figure 2: Load operation decision flow in `checkLoad()`

The load logic follows this shape:

1. **Check if tracked** — If the address is not in the pmap, it is not being tracked for lifetime purposes (e.g., it might be a global). Return early.
2. **Pointer type handling** — If the address is a pointer type:
 - If `isDeref` is `false`, this is loading the pointer *value* itself (e.g., copying a pointer), which is safe. Return.
 - If `isDeref` is `true`, this is *dereferencing* the pointer (e.g., `*ptr`). Call `checkPointerDeref()` to verify it is not invalid or null.

3. **Value type handling** — If the address is a value type, check if it is moved-from using `isValueTypeMovedFrom()`. If so, emit a diagnostic for use-after-move.

Implementation pattern from `LifetimeCheck.cpp`, simplified:

```
1 void LifetimeCheckPass::checkLoad(LoadOp loadOp) {
2     auto addr = loadOp.getAddr();
3     if (!getPmap().count(addr))
4         return; // Not tracked
5
6     // For pointer types, only check on dereference
7     if (ptrs.count(addr)) {
8         if (!loadOp.getIsDeref())
9             return; // Loading pointer value is safe
10        checkPointerDeref(addr, loadOp.getLoc());
11        return;
12    }
13
14    // For value types, check if moved-from
15    if (isValueTypeMovedFrom(addr)) {
16        checkPointerDeref(addr, loadOp.getLoc());
17        return;
18    }
19 }
```

The `isDeref` flag is set by CIRGen based on the C++ operation:

```
1 int *ptr = ...;
2 int *ptr2 = ptr; // LoadOp(ptr, isDeref=false) - loading pointer
   → value
3 int value = *ptr; // LoadOp(ptr, isDeref=true) - dereferencing
   → pointer
```

This distinction allows the pass to avoid false positives: copying a null pointer is legal, but dereferencing it is not.

2 Control flow analysis

ClangIR provides high-level control flow operations that make branch analysis easier than with raw LLVM IR. While MLIR offers a built-in dataflow analysis framework (`mlir::dataflow::DataFlowSolver`), the

LifetimeCheck pass implements custom control flow analysis for fine-grained control over state tracking and error reporting.

The pass builds on MLIR's base infrastructure, which provides the fundamental Region/Block/Operation hierarchy that structures the IR, along with traversal iterators such as `region.getBlocks()` and `block.getOperations()` for walking through the program structure. MLIR also provides operation type dispatch mechanisms (`isa<>()`, `cast<>()`) and region accessors (`getThenRegion()`, `getElseRegion()`) that enable navigation through control flow constructs. On top of this infrastructure, LifetimeCheckPass implements its own custom logic: it maintains state tracking through the pointer map (pmap) and scope management, implements state merging logic via `joinPmaps()` to handle control flow merge points, and defines the actual lifetime checking algorithms that detect use-after-move and dangling pointer errors.

The pass handles branches by merging state:

```

1 void LifetimeCheckPass::checkIf(IfOp ifOp) {
2     // Collect pmaps from all branches for joining
3     llvm::SmallVector<PMapType, 2> pmapOps;
4
5     {
6         PMapType localThenPmap = getPmap();
7         PmapGuard pmapGuard{*this, &localThenPmap};
8         checkRegionWithScope(ifOp.getThenRegion());
9         pmapOps.push_back(localThenPmap);
10    }
11
12    // In case there's no 'else' branch, use the incoming pmap
13    if (!ifOp.getElseRegion().empty()) {
14        PMapType localElsePmap = getPmap();
15        PmapGuard pmapGuard{*this, &localElsePmap};
16        checkRegionWithScope(ifOp.getElseRegion());
17        pmapOps.push_back(localElsePmap);
18    } else {
19        pmapOps.push_back(getPmap());
20    }
21
22    joinPmaps(pmapOps);
23 }

```

The merge operation conservatively handles both paths:

```

1 // If a pointer is invalid in any branch, it's invalid after the if
2 if (psetThen.count(State::getInvalid()) ||

```

```

3     psetElse.count(State::getInvalid()) {
4     mergedPset.insert(State::getInvalid());
5     }

```

Example:

```

1  std::unique_ptr<int> ptr = std::make_unique<int>(42);
2  if (condition) {
3      auto ptr2 = std::move(ptr); // ptr invalid in 'then' branch
4  } else {
5      // ptr still valid in 'else' branch
6  }
7  // After merge: ptr is potentially invalid
8  int value = *ptr; // WARNING: might be invalid

```

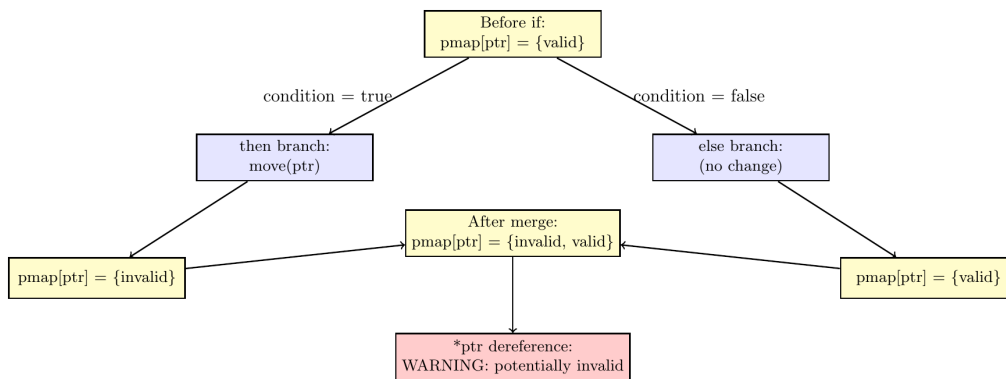


Figure 3: Control flow state merging in if/else branches

The conservative merge strategy ensures that if a variable becomes invalid in *any* branch, it is treated as potentially invalid after the merge. This prevents false negatives but may cause false positives that require user annotations to suppress.

2.1 Loop analysis

Loops present a challenge for lifetime analysis because they may execute zero or more times, and the number of iterations affects what values are valid. The LifetimeCheck pass uses a loop unrolling model that analyzes loops as if they were the first two iterations unrolled into conditional statements.

This approach, specified in P1179 §2.4.9, treats a loop:

```

1 for (/*init*/; /*cond*/; /*incr*/) {
2   /*body*/
3 }

```

as if it were:

```

1 if (/*init*/; /*cond*/) {
2   /*body*/; /*incr*/
3 }
4 if (/*cond*/) {
5   /*body*/
6 }

```

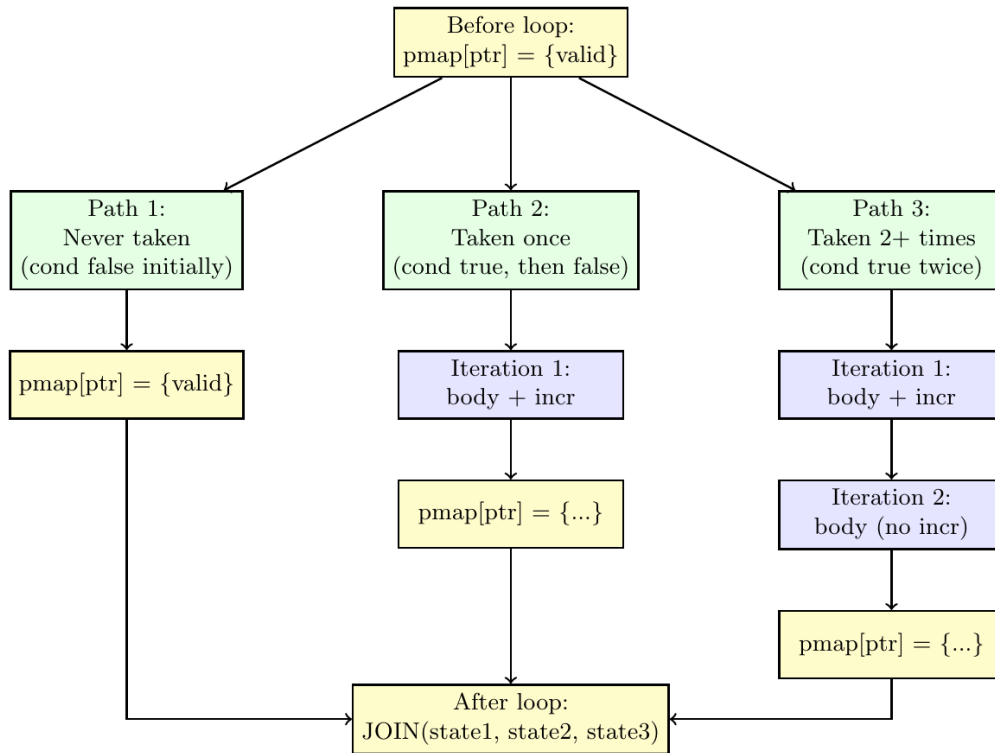


Figure 4: Loop unrolling model: three paths covering 0, 1, and 2+ iterations

The three paths represent:

1. **Never taken** — The loop condition is false from the start, so the body never executes. The pmap remains unchanged from before the loop.

2. **Taken once** — The condition is true initially, so the loop body executes once. In the implementation below, the first-taken path walks the regions in execution order but drops the step region when the loop has one. This is the shape of the pass as implemented here.
3. **Taken twice or more** — The condition is true at least twice. The pass starts from the first-taken exit pmap and checks the execution-order regions again, keeping the step region on this subsequent path.

After analyzing all three paths, the pass joins the resulting pmaps using the JOIN operation (explained below). This conservative approach ensures that:

- If a pointer becomes invalid *in any iteration*, it is treated as potentially invalid after the loop.
- If a loop may execute zero times, the analysis preserves the pre-loop state as one possibility.

Implementation pattern from `LifetimeCheck.cpp`, simplified:

```

1 void LifetimeCheckPass::checkLoop(LoopOpInterface loopOp) {
2     // Treat loop as first two iterations unrolled with if statements
3     llvm::SmallVector<PMapType, 4> pmapOps;
4     llvm::SmallVector<Region *, 4> regionsToCheck;
5
6     auto setupLoopRegionsToCheck = [&](bool isSubsequentTaken =
7     ↪ false) {
8         regionsToCheck = loopOp.getRegionsInExecutionOrder();
9         if (loopOp.maybeGetStep() && !isSubsequentTaken)
10            regionsToCheck.pop_back();
11    };
12
13    // Path 1: Never taken
14    pmapOps.push_back(getPmap());
15
16    // Path 2: Taken once (condition true, then false)
17    PMapType loopExitPmap;
18    {
19        loopExitPmap = getPmap();
20        PmapGuard pmapGuard{*this, &loopExitPmap};
21        setupLoopRegionsToCheck();
22        for (auto *region : regionsToCheck)

```

```

22     checkRegion(*region);
23     pmapOps.push_back(loopExitPmap);
24 }
25
26 // Path 3: Taken 2+ times (condition true at least twice)
27 if (getPmap() != loopExitPmap) {
28     PMapType otherTakenPmap = loopExitPmap;
29     PmapGuard pmapGuard{*this, &otherTakenPmap};
30     setupLoopRegionsToCheck(/*isSubsequentTaken=*/true);
31     for (auto *region : regionsToCheck)
32         checkRegion(*region);
33     pmapOps.push_back(otherTakenPmap);
34 }
35
36 // Conservatively merge all three paths
37 joinPmaps(pmapOps);
38 }

```

The exact region choice is an implementation detail of this pass. The important analysis property is that the never-taken pmap, the first-taken pmap, and the subsequent-taken pmap are all joined at the loop exit.

Example demonstrating loop lifetime tracking:

```

1  std::unique_ptr<int> ptr = std::make_unique<int>(42);
2  for (int i = 0; i < n; ++i) {
3      if (i == 0) {
4          auto ptr2 = std::move(ptr); // ptr invalidated on first
           ↪ iteration
5      }
6      // After loop merge: ptr potentially invalid
7  }
8  *ptr; // WARNING: pointer might be invalid

```

The analysis correctly identifies that `ptr` may be invalid after the loop because it could be moved on the first iteration (if `n > 0`).

2.2 Switch statement analysis

Switch statements are analyzed similarly to loops: they are transformed into an equivalent series of if-else statements. The pass handles switch statements in simple form—cases with single regions and explicit break or fallthrough behavior.

Per P1179 §2.4.7, a switch:

```

1  switch (a) {
2     case 1: /*1*/
3     case 2: /*2*/ break;
4     default: /*3*/
5  }

```

is treated as:

```

1  if (auto& a=a; a==1) { /*1*/ }
2  else if (a==1 || a==2) { /*2*/ }
3  else { /*3*/ }

```

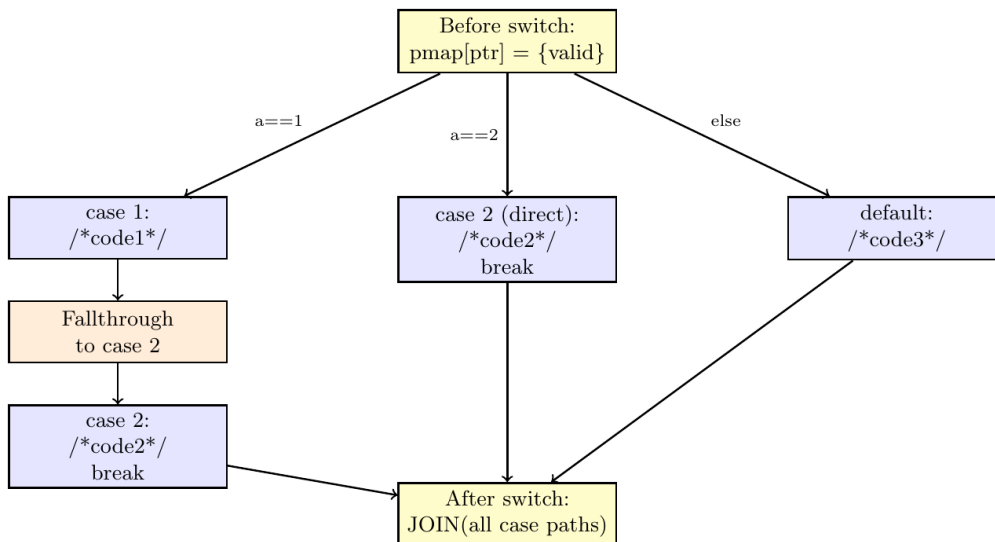


Figure 5: Switch statement analysis with fallthrough handling

Key aspects of switch analysis:

- **Fallthrough semantics** — When a case does not end with `break`, execution falls through to the next case. The pass models this by sequentially checking both case regions and merging the final state.
- **Break semantics** — When a case ends with `break`, execution jumps to after the switch. Each such path contributes independently to the final joined pmap.
- **Default case** — If present, the default case is checked as another independent path. If absent, the incoming pmap (representing “no case matched”) is included in the join.

Implementation pattern from LifetimeCheck.cpp, simplified:

```
1 void LifetimeCheckPass::checkSwitch(SwitchOp switchOp) {
2     llvm::SmallVector<PMapType, 2> pmapOps;
3
4     // Only handle switch in simple form
5     llvm::SmallVector<CaseOp> cases;
6     if (!switchOp.isSimpleForm(cases))
7         return;
8
9     auto isCaseFallthroughTerminated = [&](Region &r) -> bool {
10         Block &block = r.back();
11         auto yieldOp = dyn_cast<YieldOp>(block.back());
12         return !!yieldOp;
13     };
14
15     // Analyze each case
16     for (size_t i = 0; i < cases.size(); ++i) {
17         auto &caseOp = cases[i];
18         PMapType localPmap = getPmap();
19         PmapGuard pmapGuard{*this, &localPmap};
20
21         // Check case region
22         for (auto &region : caseOp.getRegions())
23             checkRegion(region);
24
25         // Handle fallthrough to next case
26         if (isCaseFallthroughTerminated(caseOp.getRegion(0))) {
27             if (i + 1 < cases.size()) {
28                 auto &nextCase = cases[i + 1];
29                 for (auto &region : nextCase.getRegions())
30                     checkRegion(region);
31             }
32         }
33
34         pmapOps.push_back(localPmap);
35     }
36
37     // Merge all case paths
38     joinPmaps(pmapOps);
39 }
```

Example demonstrating switch lifetime tracking:

```

1  std::unique_ptr<int> ptr = std::make_unique<int>(42);
2  switch (value) {
3      case 1:
4          auto ptr2 = std::move(ptr); // ptr invalidated in case 1
5          break;
6      case 2:
7          // ptr still valid in case 2
8          break;
9      default:
10         // ptr still valid in default
11         break;
12 }
13 // After switch: ptr potentially invalid (could have taken case 1)
14 *ptr; // WARNING: pointer might be invalid

```

2.3 The JOIN operation

The JOIN operation is the fundamental mechanism for merging program state from multiple control flow paths. It implements a conservative union strategy: if a variable has different states in different branches, the merged state must be conservative enough to cover all possibilities.

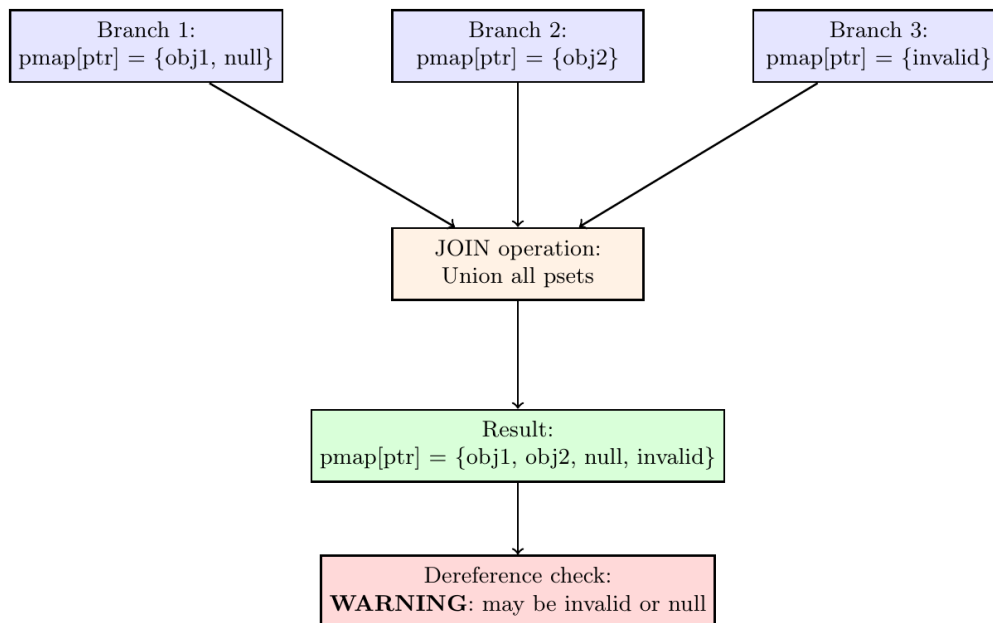


Figure 6: JOIN operation: conservative union of points-to sets from multiple branches

The JOIN algorithm (P1179 §2.3) [2] is simple but powerful:

1. For each tracked address `addr` in the `pmap`, collect its `pset` from each branch being merged.
2. Compute the *set union* of all these `psets`.
3. Replace `pmap[addr]` with the union result.

Implementation pattern from `LifetimeCheck.cpp`:

```
1 void LifetimeCheckPass::joinPmaps(SmallVectorImpl<PMapType> &pmaps)
   → {
2   for (auto &mapEntry : getPmap()) {
3     auto &val = mapEntry.first;
4
5     // Collect pset from each branch
6     PSetType joinPset;
7     for (auto &pmapOp : pmaps)
8       llvm::set_union(joinPset, pmapOp[val]);
9
10    // Update with union of all psets
11    getPmap()[val] = joinPset;
12  }
13 }
```

Why this is conservative:

- If a pointer is `{invalid}` in *any* branch, the merged `pset` contains `invalid`, so any subsequent dereference triggers a warning.
- If a pointer is `{null}` in some branches and `{valid}` in others, the merged `pset` is `{null, valid}`, indicating potential null.
- If a pointer points to different objects in different branches, the merged `pset` contains all possibilities.

Example demonstrating JOIN with three branches:

```
1 int *ptr;
2 switch (value) {
3   case 1: ptr = nullptr; break;           // pmap[ptr] = {null}
4   case 2: ptr = &x; break;               // pmap[ptr] = {x}
5   case 3: ptr = invalid_ptr; break;      // pmap[ptr] = {invalid}
6 }
7 // After JOIN: pmap[ptr] = {null, x, invalid}
8 *ptr; // WARNING: may be invalid or null
```

The conservative nature of JOIN is essential for soundness: the analysis would rather emit false positives (warnings for safe code) than false negatives (missing real bugs). This aligns with the C++ Core Guidelines [1] philosophy that lifetime safety analysis should be conservative but suppressible.

3 Scope and lifetime management

One of the fundamental operations in the C++ Core Guidelines lifetime safety profile (P1179) is the KILL operation. Understanding KILL is essential to understanding how the LifetimeCheck pass prevents use-after-free bugs.

3.1 Lexical scope tracking

The pass tracks lexical scopes using a `LexicalScopeContext` stack. Each scope maintains a set of local values declared within it. When a scope ends (e.g., closing brace of a function or block), all local values in that scope must be destroyed, and any pointers referencing them must be invalidated.

This is managed using the RAII pattern with `LexicalScopeGuard`:

```
1 void LifetimeCheckPass::checkRegionWithScope(Region &region) {
2     // Create new scope
3     LexicalScopeContext lexScope{currFunc};
4     LexicalScopeGuard scopeGuard{*this, &lexScope};
5
6     // Check operations in this scope
7     for (auto &block : region)
8         checkBlock(block);
9
10    // Scope guard destructor runs here, calling kill() for all locals
11 }
```

When the `LexicalScopeGuard` destructor runs (at scope exit), it calls `kill()` for each local value that is going out of scope.

3.2 The KILL operation

When a local `Owner` goes out of scope, the lifetime analysis must invalidate all pointers that reference it or the resource it owns. This is the KILL operation from the C++ Core Guidelines lifetime safety profile [2]:

KILL(x) means to replace all occurrences of x and x' and x'' (etc.) in the pmap with invalid.

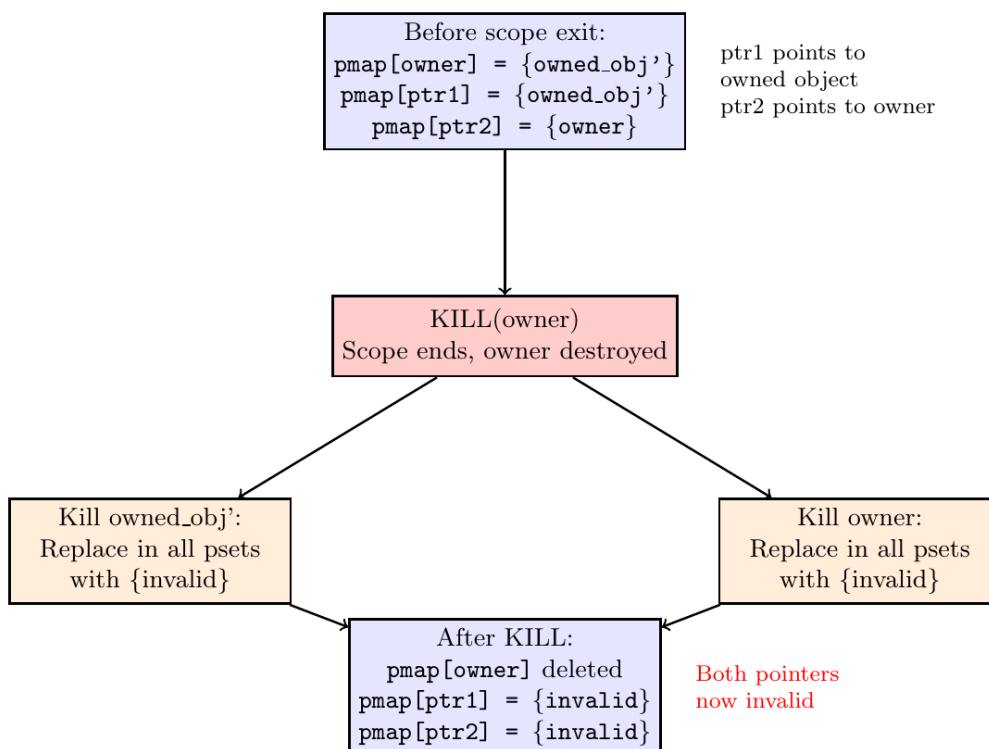


Figure 7: KILL operation cascading invalidation when owner goes out of scope

The KILL operation has cascading semantics:

- When `owner` is killed, both `owner` itself and the owned object `owner'` are invalidated.
- All pointers in the `pmap` that reference `owner` or `owner'` have their psets updated to `{invalid}`.
- The `owner` is removed from the category tracking sets when its scope ends. Other `pmap` entries that mention `owner` or `owner'` are rewritten to `{invalid}`.

Implementation pattern from `LifetimeCheck.cpp`, simplified:

```

1 void LifetimeCheckPass::killInPset(mlir::Value ptrKey,
2                                   const State &s,
3                                   InvalidStyle invalidStyle,
4                                   mlir::Location loc) {

```

```

5     auto &pset = getPmap()[ptrKey];
6     if (pset.contains(s)) {
7         pset.erase(s);
8         markPsetInvalid(ptrKey, invalidStyle, loc);
9     }
10 }
11
12 // KILL(x): replace all occurrences of x, x', x'' in pmap with
13 → invalid
14 void LifetimeCheckPass::kill(const State &s,
15                               InvalidStyle invalidStyle,
16                               mlir::Location loc) {
17     assert(s.hasValue() && "does not know how to kill other types");
18     mlir::Value v = s.getData();
19
20     for (auto &mapEntry : getPmap()) {
21         auto ptr = mapEntry.first;
22
23         // Skip the entry being deleted
24         if (v == ptr)
25             continue;
26
27         // Replace all occurrences of x' (owned object)
28         if (s.isLocalValue() && owners.count(v))
29             killInPset(ptr, State::getOwnedBy(v), invalidStyle, loc);
30
31         // Replace all occurrences of x (the value itself)
32         killInPset(ptr, s, invalidStyle, loc);
33     }
34
35     // Remove scoped local from category tracking sets
36     if (invalidStyle == InvalidStyle::EndOfScope) {
37         owners.erase(v);
38         ptrs.erase(v);
39         tasks.erase(v);
40     }
41 }
42
43 // LexicalScopeGuard destructor triggers KILL for all locals
44 void LifetimeCheckPass::LexicalScopeGuard::cleanup() {
45     auto *localScope = Pass.currScope;
46     for (auto pointee : localScope->localValues)
47         Pass.kill(State::getLocalValue(pointee),

```

```

47         InvalidStyle::EndOfScope,
48         getEndLocForHist(*localScope));
49     }

```

Example demonstrating KILL in action:

```

1  void example() {
2      std::unique_ptr<int> owner = std::make_unique<int>(42);
3      int *ptr1 = owner.get();    // ptr1 -> owned object
4      auto *ptr2 = &owner;       // ptr2 -> owner itself
5
6      {
7          int *ptr3 = owner.get(); // ptr3 -> owned object
8          // End of inner scope: KILL(ptr3)
9          // pmap[ptr3] deleted
10     }
11     // ptr3 no longer exists
12
13     *ptr1; // Still OK here - owner still alive
14
15     // End of function scope: KILL(owner)
16     // This invalidates both ptr1 and ptr2
17     // because one points to owner' and one points to owner
18 }
19
20 void dangling_reference() {
21     int *dangling;
22     {
23         std::vector<int> vec = {1, 2, 3};
24         dangling = vec.data();
25         // End of scope: KILL(vec)
26         // pmap[dangling] = {invalid} (points to destroyed resource)
27     }
28     *dangling; // ERROR: use of invalid pointer
29 }

```

The KILL operation is fundamental to preventing use-after-free bugs because it ensures that when an object is destroyed, all pointers to it are marked invalid. Any subsequent dereference of these pointers will trigger a diagnostic from `checkPointerDeref()`.

The relationship between KILL and the three type categories is important:

- **Owners** are killed when they go out of scope, which invalidates all pointers to the owned resource.

- **Pointers** are killed when they go out of scope, but this only affects pointers-to-pointers (rare).
- **Values** are killed when they go out of scope, invalidating any pointers directly to the value’s address (e.g., `&x` where `x` is an int).

This is why the pass must carefully track which category each variable belongs to—the semantics of KILL depend on it.

4 Coroutines and async code

Coroutines introduce unique lifetime challenges because they can suspend execution and resume later, potentially after local variables have gone out of scope. The `LifetimeCheck` pass has specialized handling for coroutine operations to detect these bugs.

4.1 Await operation handling

The `co_await` operator suspends the active coroutine and may execute several regions of code: the awaiter’s `await_ready()`, `await_suspend()`, and `await_resume()` methods. The pass conservatively analyzes all regions sequentially and joins their resulting states.

Implementation pattern from `LifetimeCheck.cpp`:

```

1 void LifetimeCheckPass::checkAwait(AwaitOp awaitOp) {
2     // Conservative: assume all regions execute sequentially
3     llvm::SmallVector<PMapType, 4> pmapOps;
4
5     for (auto r : awaitOp.getRegions()) {
6         PMapType regionPmap = getPmap();
7         PmapGuard pmapGuard{*this, &regionPmap};
8         checkRegion(*r);
9         pmapOps.push_back(regionPmap);
10    }
11
12    // Join states from all awaiter methods
13    joinPmaps(pmapOps);
14 }

```

This conservative approach ensures that if any awaiter method invalidates a pointer, the merged state reflects that possibility. Example:

```
1 Task<void> example(std::unique_ptr<int> ptr) {
2     co_await suspendPoint();
3     // After await: ptr might still be valid (conservative)
4     *ptr; // OK if suspendPoint doesn't move ptr
5 }
```

4.2 Coroutine task lifetime tracking

Coroutine tasks (the objects returned by coroutine functions) require special lifetime tracking because they can capture references to local variables passed as arguments. When a task is created, the pass tracks which locals are bound to it, ensuring that when those locals are destroyed, the task is invalidated.

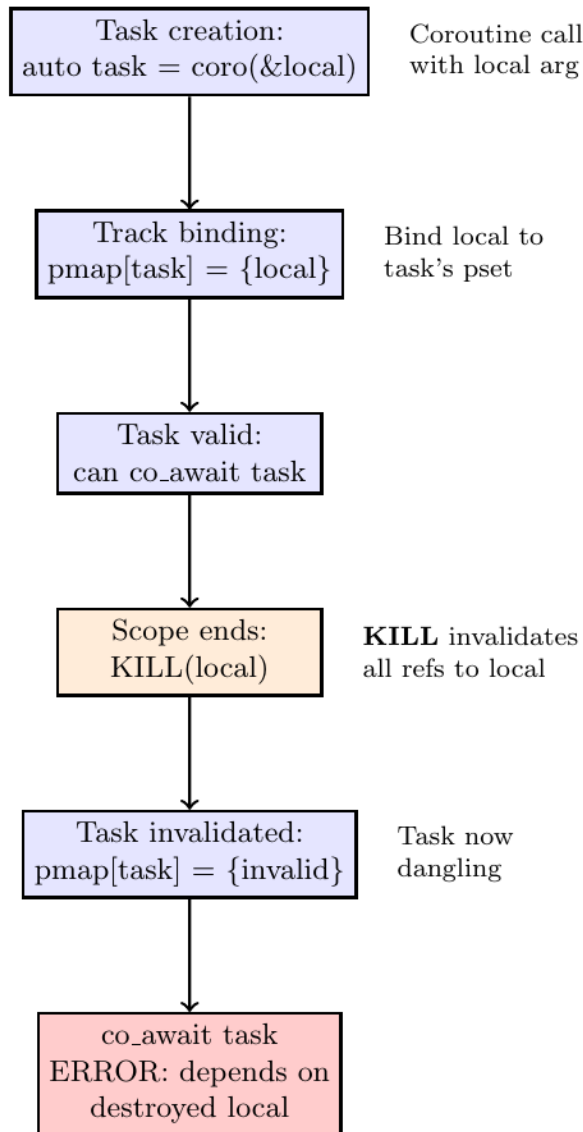


Figure 8: Coroutine task lifetime: tracking captured local references

The task tracking logic is implemented in two parts:

1. Identifying coroutine calls:

```

1 void LifetimeCheckPass::trackCallToCoroutine(CallOp callOp) {
2   if (auto calleeFuncOp = callOp.getDirectCallee(theModule)) {
3     if (calleeFuncOp.getCoroutine() ||
4         (calleeFuncOp.isDeclaration() && callOp->getNumResults() >
5         ↪ 0 &&

```

```

5         isTaskType(callOp->getResult(0))) {
6         currScope->localTempTasks.insert(callOp->getResult(0));
7     }
8     return;
9 }
10 // Handle indirect calls to coroutines, for instance when
11 // lambda coroutines are involved with invokers.
12 if (callOp->getNumResults() > 0 &&
13     ↪ isTaskType(callOp->getResult(0))) {
14     currScope->localTempTasks.insert(callOp->getResult(0));
15 }

```

2. Binding locals to task pset (covered earlier in `checkCoroTaskStore()`):

When the task is stored to a variable, any local values passed as arguments are added to the task's pset. When those locals are destroyed (KILL operation), the task becomes invalid.

Common coroutine lifetime bug example:

```

1 Task<int> dangling_task_example() {
2     int local = 42;
3     auto task = async_compute(&local); // task captures &local
4     // End of scope: KILL(local)
5     // pmap[task] contains local, so task is invalidated too
6     return task; // WARNING: task depends on destroyed local
7 }
8
9 Task<void> use_after_await() {
10    Task<int> task = compute();
11    co_await std::move(task);
12    // task is now moved-from
13    int result = co_await task; // ERROR: use-after-move
14 }

```

4.3 Temporary vs persistent tasks

Most temporaries in C++ can be ignored for move tracking because they are destroyed at the end of the full expression. However, coroutine task temporaries are an exception because they may be captured by coroutine frames and outlive their expression.

```

1 bool LifetimeCheckPass::isSkippableTemporary(mlir::Value v) {
2     auto allocaOp = v.getDefiningOp<cir::AllocaOp>();

```

```

3   if (!allocaOp)
4       return false;
5
6   // Temporaries have "ref.tmp" prefix
7   auto name = allocaOp.getName();
8   if (!name.starts_with("ref.tmp"))
9       return false;
10
11  // IMPORTANT: Do not skip coroutine tasks
12  // They need lifetime tracking even as temporaries
13  if (isTaskType(v))
14      return false;
15
16  return true; // Other temporaries can be skipped
17 }

```

This distinction allows:

```

1 // Regular temporary - can skip move tracking
2 int x = std::move(SomeClass()); // temporary SomeClass() can be
   → skipped
3
4 // Task temporary - must track
5 auto result = co_await createTask(); // task temporary is tracked

```

4.4 Lambda captures by reference

Lambdas that capture local variables by reference create similar lifetime issues to coroutine tasks. The `checkLambdaCaptureStore()` function (covered earlier) binds captured locals to the lambda's pset.

Common lambda capture bug:

```

1 std::function<int()> create_lambda() {
2     int local = 42;
3     auto lambda = [&local]() { return local + 1; };
4     // lambda captures local by reference
5     // End of scope: KILL(local) invalidates lambda
6     return lambda; // WARNING: lambda captures destroyed local
7 }
8
9 void lambda_use_after_move() {
10    std::vector<int> vec = {1, 2, 3};
11    auto lambda = [vec = std::move(vec)]() { return vec.size(); };

```

```

12     // vec is moved-from
13     auto size = vec.size(); // ERROR: use-after-move
14 }

```

The pass detects these patterns by tracking stores to lambda capture fields and binding the captured locals to the lambda's pset. When the locals go out of scope, the KILL operation propagates through the pset and invalidates the lambda.

5 Return value safety

Returning references or pointers to local variables is a classic C++ bug. The LifetimeCheck pass handles one important form of this problem: returning lambdas that may capture local references. Plain reference and pointer returns are a natural extension point.

Implementation pattern from LifetimeCheck.cpp:

```

1 void LifetimeCheckPass::checkReturn(ReturnOp retOp) {
2     // Invalidate all local values on return
3     if (retOp.getNumOperands() == 0)
4         return;
5
6     auto retTy = retOp.getOperand(0).getType();
7     // Currently only handles lambda return types
8     if (!isLambdaType(retTy))
9         return;
10
11     // The return value is loaded from the return slot
12     auto loadOp = retOp.getOperand(0).getDefiningOp<cir::LoadOp>();
13     assert(loadOp && "expected cir.load");
14     if (!loadOp.getAddr().getDefiningOp<cir::AllocaOp>())
15         return;
16
17     // Track lambda for later checking
18     // Actual check happens at scope exit (LexicalScopeGuard)
19     currScope->localRetLambdas.insert(
20         std::make_pair(loadOp.getAddr(), loadOp.getLoc()));
21 }

```

The clever aspect is that the pass does not check the lambda immediately. Instead, it defers the check until the scope ends (in `LexicalScopeGuard::cleanup()`). This allows it to determine which locals the lambda captured and whether any are being destroyed at scope exit.

From `LexicalScopeGuard::cleanup()`:

```
1 void LifetimeCheckPass::LexicalScopeGuard::cleanup() {
2     auto *localScope = Pass.currScope;
3     // Kill all local values going out of scope
4     for (auto pointee : localScope->localValues)
5         Pass.kill(State::getLocalValue(pointee),
6                 ↪ InvalidStyle::EndOfScope,
7                 getEndLocForHist(*localScope));
8
9     // Check returned lambdas for dangling references
10    for (auto l : localScope->localRetLambdas)
11        Pass.checkPointerDeref(l.first, l.second,
12                               ↪ DereferStyle::RetLambda);
13 }
14 }
```

Example demonstrating return safety checking:

```
1 auto return_dangling_lambda() {
2     int local = 42;
3     auto lambda = [&local]() { return local; };
4     return lambda; // WARNING: lambda captures local by reference
5     // At scope exit:
6     // 1. KILL(local) invalidates all refs to local
7     // 2. checkPointerDeref on lambda detects it's invalid
8 }
9
10 int& return_reference_to_local() {
11     int local = 42;
12     return local; // BUG: returning reference to local
13     // Future work for this pass: diagnose plain reference returns
14     ↪ too.
15 }
```

The implementation focuses on lambda returns because they are common in modern C++ (especially with `std::function` and callbacks). One natural extension would be to detect all dangling reference returns, not just lambdas.

6 Implementation patterns and best practices

During implementation, several patterns emerged that improve code quality and maintainability.

6.1 Semantic helper methods

Instead of complex inline conditions, extract logic into methods with semantic names:

```
1 // Bad: hard to understand at a glance
2 if (getPmap().count(srcAddr) &&
3     getPmap()[srcAddr].count(State::getInvalid()) &&
4     !owners.count(srcAddr) &&
5     !ptrs.count(srcAddr)) {
6     // ...
7 }
8
9 // Good: clear intent
10 if (isValueTypeMovedFrom(srcAddr)) {
11     // ...
12 }
```

6.2 Composable building blocks

Build complex checks from simple, reusable pieces:

```
1 // Atomic checks
2 bool isValueType(mlir::Value addr);
3 bool hasInvalidState(mlir::Value addr);
4
5 // Composed checks
6 bool isValueTypeMovedFrom(mlir::Value addr) {
7     return isValueType(addr) && hasInvalidState(addr);
8 }
```

6.3 API symmetry

Maintain parallel naming for related operations:

```
1 // For owner types
2 void markOwnerAsMovedFrom(mlir::Value addr, mlir::Location loc);
3
4 // For pointer/value types (added for symmetry)
5 void markPointerOrValueTypeAsMovedFrom(mlir::Value addr,
6                                         mlir::Location loc);
```

6.4 Early-exit pattern

Reduce nesting with early returns:

```
1 // Before: deeply nested
2 if (auto allocaOp = getDefiningOp<AllocaOp>()) {
3     if (isValueType(addr)) {
4         if (!hasInvalidState(addr)) {
5             // actual logic buried here
6         }
7     }
8 }
9
10 // After: early exits
11 auto allocaOp = getDefiningOp<AllocaOp>();
12 if (!allocaOp)
13     return;
14 if (!isValueType(addr))
15     return;
16 if (hasInvalidState(addr))
17     return;
18 // actual logic at top level
```

Refactoring pattern	Before	After	Benefit
Semantic helper methods	4-line condition	1-line call	Readability
Duplicate code elimination	8+ locations	1 method	Maintainability
Early-exit pattern	3-level nesting	Flat structure	Clarity
API symmetry	Inconsistent names	Parallel naming	Discoverability
Composable helpers	Monolithic checks	Building blocks	Reusability

Figure 9: Refactoring patterns and their impact

These patterns are small, but they matter for an analysis pass. Most checks are not complicated individually; the difficulty is keeping many small cases consistent. Semantic helper methods and symmetric APIs make it easier to review the pass and to extend it later.

7 Limitations and future work

The LifetimeCheck pass has known limitations:

- **Interprocedural analysis:** The pass analyzes each function independently. It does not track lifetime across function boundaries, which can lead to false negatives when a dangling pointer is returned from a function.
- **Temporary detection:** The pass uses a string prefix (`ref.tmp`) to identify temporaries. A more robust approach would be adding an `is_temporary` attribute to `AllocaOp` during CIRGen.
- **Field-sensitive analysis:** The pass tracks first-level aggregate fields that are actually reached through `GetMemberOp` uses. Nested aggregate fields and unused fields are not fully modeled.
- **Custom smart pointers:** The special smart-pointer path recognizes only `unique_ptr` and `shared_ptr` names. User-defined smart pointers following the same patterns are not handled specially.

Future improvements could address these by:

- Adding summary information to function signatures.
- Enhancing CIRGen to mark temporaries explicitly.
- Implementing field-sensitive tracking using CIR's `GetMemberOp`.
- Allowing user annotations for custom smart pointer types.

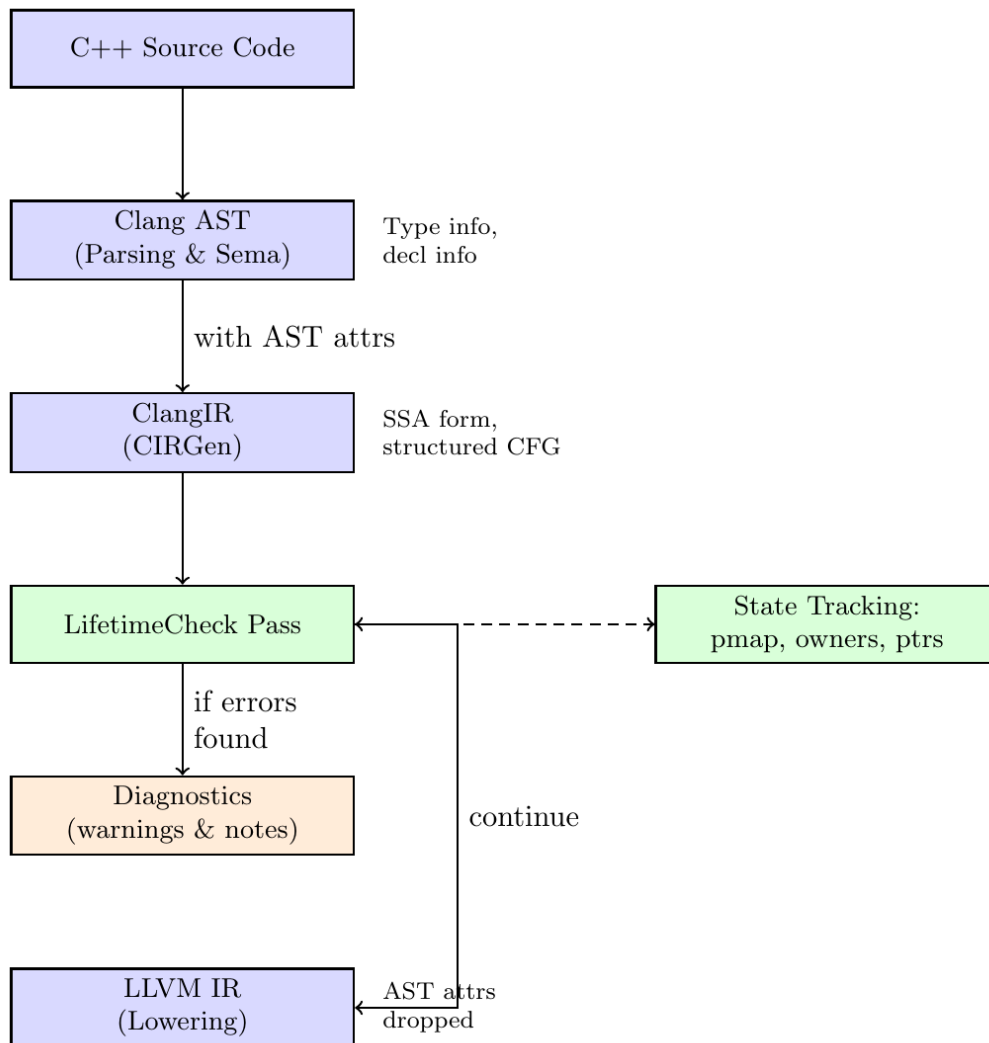


Figure 10: Workflow from C++ source to lifetime diagnostics

The LifetimeCheck pass operates on ClangIR after CIRGen and before lowering to LLVM IR. This is the useful point in the pipeline: the pass still has access to AST semantic information, but it can also use IR structure for data-flow reasoning.

8 Lessons learned

Studying the LifetimeCheck pass demonstrates several advantages of ClangIR for compiler analysis:

1. **AST attributes preserve semantics:** Access to AST information through interfaces makes it easy to answer questions like “Is this parameter an rvalue reference?” or “Is this type a `std::unique_ptr`?” that would be difficult or impossible with LLVM IR.
2. **High-level operations clarify intent:** CIR’s `StoreOp`, `LoadOp`, and `CallOp` carry more semantic meaning than LLVM’s generic instructions, making the analysis code more readable and maintainable.
3. **SSA simplifies dataflow:** Unlike AST where variables can be re-assigned, CIR’s SSA form makes it easy to track value flow. Each `mlir::Value` has exactly one definition.
4. **Structured control flow:** CIR’s `IfOp`, `SwitchOp`, and `LoopOp` are easier to analyze than LLVM’s unstructured basic block CFG. Branch merging logic is straightforward.
5. **Location tracking:** MLIR requires every operation to have a source location, making diagnostic emission trivial compared to LLVM IR where location metadata is optional.

The `LifetimeCheck` pass serves as a practical template for building CIR analysis passes. The patterns demonstrated—AST attribute access, state tracking, operation visitation, and diagnostic emission—apply broadly to many kinds of static analysis tasks.

References

- [1] Bjarne Stroustrup and Herb Sutter. C++ Core Guidelines. <https://isocpp.org/guidelines>, 2025. Lifetime Safety Profile: <https://isocpp.github.io/CppCoreGuidelines/CppCoreGuidelines#SS-lifetime>.
- [2] Herb Sutter. Lifetime safety: Preventing common dangling. C++ Standards Committee Paper P1179R1, ISO/IEC JTC1/SC22/WG21, November 2019.