

AST Semantics and Use-After-Move Detection

Ivan Murashko

Contents

1 Leveraging AST semantics	1
1.1 Detecting move operations	2
1.2 Smart pointer detection	4
1.3 Smart pointer special handling	5
2 Detecting use-after-move	8
2.1 Marking values as moved-from	8
2.2 Use-after-move detection algorithm	10
3 Conclusion	12

This article continues from [Building the LifetimeCheck Pass](#). The previous part defined the state model from the lifetime safety profile [1]. Here we focus on the semantic information that makes this model useful for C++: AST attributes, smart pointer recognition, and move operations.

1 Leveraging AST semantics

ClangIR operations carry AST attributes that preserve semantic information. The LifetimeCheck pass uses these through the AST attribute interfaces defined in `ASTAttrInterfaces.td`. This is one of the most important parts of the example: semantic questions are wrapped in named interface methods instead of being scattered as raw AST operations throughout the pass.

To understand how this works, let's see what CIR looks like for a simple program:

	ClangIR (simplified):
	1 cir.func @test() {
	2 %ptr = cir.alloca
	→ !cir.ptr<!ty_unique_ptr>
	3 loc("test.cpp":2:3)
	4 %val = cir.call @make_unique()
	5 loc("test.cpp":3:9)
C++ source:	6 cir.store %val, %ptr
1 void test() {	7 loc("test.cpp":3:7)
2 std::unique_ptr<int> ptr;	8 %ptr2 = cir.alloca
3 ptr = std::make_unique<int>	→ !cir.ptr<!ty_unique_ptr>
↪ (42);	9 loc("test.cpp":4:8)
4 auto ptr2 = std::move(ptr);	10 %moved = cir.call @move(%ptr)
5 }	11 loc("test.cpp":4:15)
	12 cir.store %moved, %ptr2
	13 loc("test.cpp":4:8)
	14 cir.return
	15 }

Figure 1: C++ code and the corresponding ClangIR representation

Key observations:

- Each CIR operation has a source location attached to it.
- Types are preserved (!ty_unique_ptr).
- Operations are in SSA form (%ptr, %val, etc.).
- High-level operations (alloca, call, store) match C++ semantics.

1.1 Detecting move operations

To detect if a function parameter is an rvalue reference (T&&), indicating a potential move, the pass uses the following implementation. This demonstrates the key pattern of *type categorization* in move semantics. The pass handles Owner, Pointer, and Value types differently because they have distinct moved-from semantics:

```

1 void LifetimeCheckPass::checkArgForRValueRef(
2     CallOp callOp, unsigned argIdx,
3     ASTFunctionDeclInterface funcDecl) {
4     // Use interface method instead of direct AST access
5     if (!funcDecl.isParamRValueReference(argIdx))
6         return;
7

```

```

8   auto arg = callOp.getArgOperand(argIdx);
9
10  // LoadOp means by-value, not a move - check for moved-from value
11  if (auto loadOp = arg.getDefiningOp<cir::LoadOp>()) {
12      auto srcAddr = loadOp.getAddr();
13      if (isValueTypeMovedFrom(srcAddr)) {
14          checkPointerDeref(srcAddr, callOp.getLoc());
15      }
16      return;
17  }
18
19  // Handle move semantics for address arguments
20  mlir::Value addr = arg;
21  if (!addr.getDefiningOp<cir::AllocaOp>() ||
22      → !getPmap().count(addr))
23      return;
24
25  // Type categorization: Owner vs Pointer vs Value
26  // Each category has different moved-from semantics
27
28  if (owners.count(addr)) {
29      // Owner types carrying [[gsl::Owner]]
30      if (getPmap()[addr].count(State::getInvalid()) ||
31          getPmap()[addr].count(State::getNullPtr())) {
32          checkPointerDeref(addr, callOp.getLoc());
33          return;
34      }
35      if (!isSkippableTemporary(addr))
36          markOwnerAsMovedFrom(addr, callOp.getLoc());
37      return;
38  }
39
40  if (ptrs.count(addr)) {
41      // Pointer types (raw pointers, references, iterators)
42      if (getPmap()[addr].count(State::getInvalid())) {
43          checkPointerDeref(addr, callOp.getLoc());
44          return;
45      }
46      if (!isSkippableTemporary(addr))
47          markPointerOrValueTypeAsMovedFrom(addr, callOp.getLoc());
48      return;
49  }

```

```

50 // Value types (primitives, structs without pointer semantics)
51 if (getPmap()[addr].count(State::getInvalid())) {
52     checkPointerDeref(addr, callOp.getLoc());
53     return;
54 }
55 if (!isSkippableTemporary(addr))
56     markPointerOrValueTypeAsMovedFrom(addr, callOp.getLoc());
57 }

```

The Pointer branch describes a conservative implementation. It is not the same rule as ordinary raw pointer copy: copying a raw pointer preserves the points-to set, while this rvalue-reference path marks the tracked Pointer-category address moved-from after the invalid-state check.

The interface method `isParamRValueReference` is defined in the AST attribute interface TableGen file:

```

1 InterfaceMethod<"", "bool", "isParamRValueReference",
2             (ins "unsigned":$paramIdx), [{}],
3     /*defaultImplementation=*/ [{
4         if (paramIdx >= $_attr.getAst()->getNumParams())
5             return false;
6         auto *param = $_attr.getAst()->getParamDecl(paramIdx);
7         return param->getType()->isRValueReferenceType();
8     }]
9 >

```

This approach is better than scattering direct AST access throughout the pass because:

- Interface methods provide a local API for the analysis.
- They handle null checks and edge cases.
- They are reusable across multiple passes.
- They clearly document what AST information is being accessed.

1.2 Smart pointer detection

Standard library smart pointers have special semantics: after a move, they are guaranteed to be null (unlike general owner types). The pass uses an AST interface from the implementation to detect them:

```

1 // In ASTAttrInterfaces.td
2 InterfaceMethod<"", "bool", "isSmartPointerOwner", (ins), [{}],
3   /*defaultImplementation=*/ [{
4     if (!$_attr.getAst()->getDeclContext()->isStdNamespace())
5       return false;
6     llvm::StringRef name = $_attr.getAst()->getName();
7     return name == "unique_ptr" || name == "shared_ptr";
8   }]
9 >

```

Usage in the pass:

```

1 bool isSmartPointer = false;
2 if (auto recordTy = mlir::dyn_cast<cir::RecordType>(type)) {
3   if (auto astAttr = recordTy.getAst()) {
4     isSmartPointer = astAttr.isSmartPointerOwner();
5   }
6 }
7
8 if (isSmartPointer) {
9   // Some smart pointer methods can be called even after move:
10  // get/reset/operator bool for unique_ptr and shared_ptr,
11  // release for unique_ptr only.
12 }

```

1.3 Smart pointer special handling

Smart pointers (`std::unique_ptr`, `std::shared_ptr`) have a well-defined empty state after move in the model used by the ClangIR tests. This allows certain methods to be called safely because they inspect or reset the object without dereferencing the stored pointer:

```

1 bool LifetimeCheckPass::isSmartPointerSafeMethod(
2   llvm::StringRef methodName) {
3   // These methods are safe to call on moved-from smart pointers
4   // This helper is name-based; release is unique_ptr-only.
5   return methodName == "get" ||
6         methodName == "release" ||
7         methodName == "reset" ||
8         methodName == "operator bool";
9 }

```

It is worth noting that this helper is intentionally simple. It checks the method name, not the precise smart-pointer kind, so `release()` must be read as a `std::unique_ptr` case.¹

The pass checks method calls on potentially invalid owners:

```

1 void LifetimeCheckPass::checkOperators(
2     CallOp callOp, ASTCXXMethodDeclInterface m) {
3     auto addr = getThisParamOwnerCategory(callOp);
4     if (!addr)
5         return;
6
7     // Check if 'this' is in invalid state
8     if (!getPmap()[addr].count(State::getInvalid()))
9         return;
10
11    // Special handling for smart pointers
12    bool isSafeMethod = false;
13    if (isSmartPointerType(addr.getType(), IsSmartPointerTyCache)) {
14        std::string methodName = m.getDeclName().getAsString();
15        isSafeMethod = isSmartPointerSafeMethod(methodName);
16    }
17
18    if (!isSafeMethod) {
19        // Unsafe operation on moved-from object
20        checkPointerDeref(addr, callOp.getLoc());
21    }
22 }

```

This allows code like:

```

1 std::unique_ptr<int> ptr = std::make_unique<int>(42);
2 std::unique_ptr<int> ptr2 = std::move(ptr);
3 if (ptr) { // OK: operator bool is safe
4     // ...
5 }
6 ptr.reset(); // OK: reset is safe
7 int *p = ptr.get(); // OK: get is safe
8 int value = *ptr; // ERROR: operator* requires non-null!

```

¹The implementation comment for `isSmartPointerSafeMethod()` is broader than the modeled standard-library API: `release()` is a `std::unique_ptr` operation, not a `std::shared_ptr` operation. The helper can be made more precise by making safe-method classification aware of the smart-pointer kind, or by keeping the comment and tests explicitly `std::unique_ptr`-only for `release()`.

Reinitializing a moved-from smart pointer with `reset(new int(...))` is valid C++, but the tests for this pass do not model that reinitialization path. The example above uses `reset()` because it leaves the pointer empty, so the subsequent dereference is genuinely invalid.

The C++ standard specifies these requirements. For `operator*` and `operator->`:

```
[unique.ptr.single.observers]
[util.smartptr.shared.obs]
“operator* and operator-> require that get() != nullptr”
```

Method	Applies to	OK?	Behavior on moved-from
Safe methods (allowed on moved-from smart pointers)			
<code>get()</code>	unique/shared	Yes	Returns the stored pointer, possibly null
<code>release()</code>	unique only	Yes	Releases ownership and returns the pointer
<code>reset()</code>	unique/shared	Yes	Clears or replaces the stored pointer
<code>operator bool</code>	unique/shared	Yes	Tests whether the stored pointer is non-null
Unsafe methods (require non-null pointer)			
<code>operator*</code>	unique/shared	No	Requires a non-null stored pointer
<code>operator-></code>	unique/shared	No	Requires a non-null stored pointer

Figure 2: Smart pointer methods: safe vs unsafe on moved-from instances

This distinction is crucial: the pass allows checking a moved-from smart pointer with `if (ptr)` but reports an error for `*ptr`.

2 Detecting use-after-move




Figure 3: State transitions for tracked values in LifetimeCheck

One of the pass's key responsibilities is detecting use-after-move bugs. Let's trace through the algorithm.

2.1 Marking values as moved-from

When the pass encounters a move operation, it marks the source as invalid:

```
1 void LifetimeCheckPass::checkMovedFromValue(StoreOp storeOp) {
2     auto loadOp = storeOp.getValue().getDefiningOp<cir::LoadOp>();
3     if (!loadOp)
4         return;
5
6     auto srcAddr = loadOp.getAddr();
7
8     // 1. Check if source is already moved-from (use-after-move!)
9     if (isValueTypeMovedFrom(srcAddr)) {
10        checkPointerDeref(srcAddr, storeOp.getLoc());
11        return;
12    }
13
14    // 2. Track rvalue initialization (e.g., int b(std::move(a)))
15    auto destAddr = storeOp.getAddr();
16    auto allocaOp = destAddr.getDefiningOp<cir::AllocaOp>();
17    if (!allocaOp)
18        return;
19
20    // Note: Actual implementation also validates
21    //   ↪ getPmap().count(srcAddr)
22    if (isValueType(destAddr) && isValueType(srcAddr)) {
```

```

22     if (!hasInvalidState(srcAddr) && !loadOp.getIsDeref()) {
23         // Mark source as moved-from
24         markPointerOrValueTypeAsMovedFrom(srcAddr, storeOp.getLoc());
25     }
26 }
27 }

```

Helper methods make the logic clear:

```

1  bool LifetimeCheckPass::isValueType(mlir::Value addr) {
2      return !owners.count(addr) && !ptrs.count(addr);
3  }
4
5  bool LifetimeCheckPass::hasInvalidState(mlir::Value addr) {
6      return getPmap().count(addr) &&
7             getPmap()[addr].count(State::getInvalid());
8  }
9
10 bool LifetimeCheckPass::isValueTypeMovedFrom(mlir::Value addr) {
11     return isValueType(addr) && hasInvalidState(addr);
12 }

```

Note the use of semantic helper methods with descriptive names. This pattern keeps conditionals readable and avoids complex inline expressions.

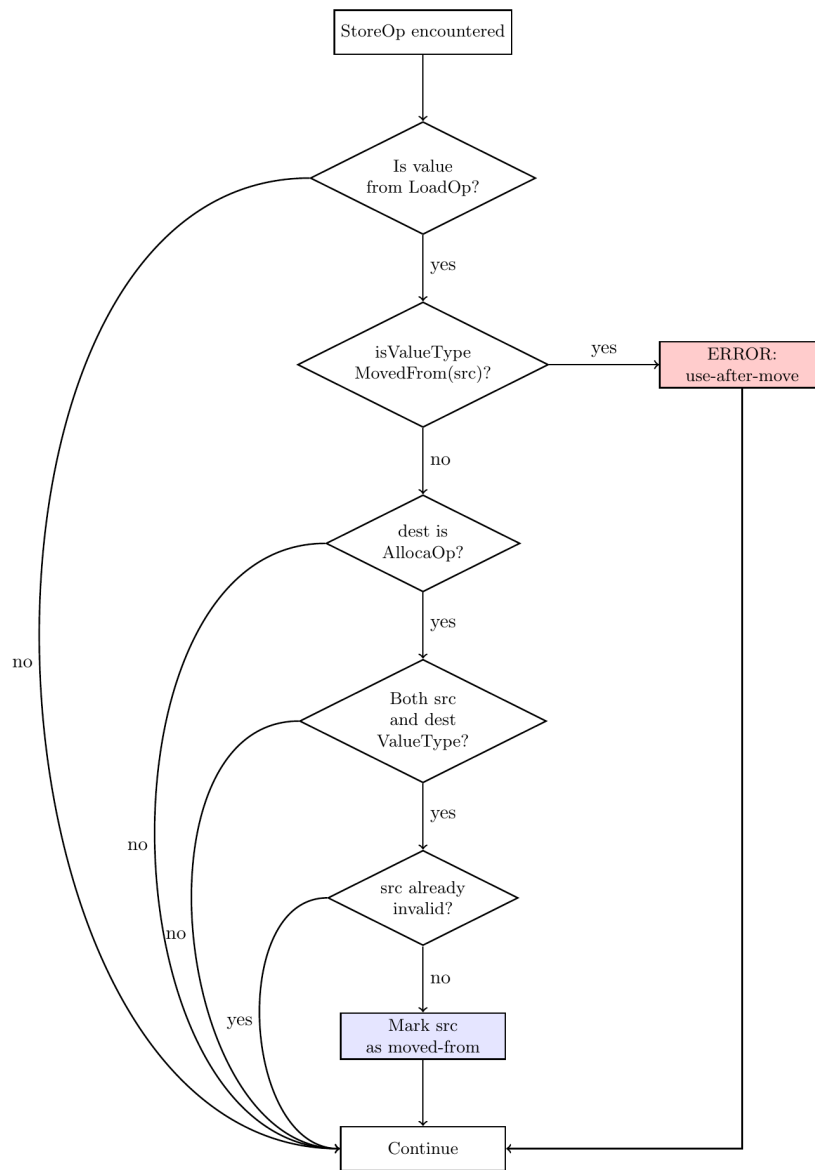


Figure 4: Algorithm flow for `checkMovedFromValue()`

The flowchart in Figure 4 shows two roles of `checkMovedFromValue()`. It detects use-after-move errors and tracks state changes by marking sources as moved-from during initialization.

2.2 Use-after-move detection algorithm

When loading from a moved-from value:

```

1 void LifetimeCheckPass::checkLoad(LoadOp loadOp) {
2     auto addr = loadOp.getAddr();
3     if (!getPmap().count(addr))
4         return;
5
6     // For pointer types, only check on dereference
7     if (ptrs.count(addr)) {
8         if (!loadOp.getIsDeref())
9             return;
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 checkPointerDeref method emits a diagnostic:

```

1 void LifetimeCheckPass::checkPointerDeref(mlir::Value addr,
2                                           mlir::Location loc,
3                                           DerefStyle derefStyle) {
4     // ... validation checks ...
5
6     auto varName = getVarNameFromValue(addr);
7     auto D = emitWarning(loc);
8
9     bool isValueType = this->isValueType(addr);
10
11    if (isValueType)
12        D << "use of moved-from value '" << varName << "'";
13    else
14        D << "use of invalid pointer '" << varName << "'";
15
16    // Emit history showing where it became invalid
17    emitInvalidHistory(D, addr, loc, derefStyle);
18 }

```

Example diagnostic output:

```

1 warning: use of moved-from value 'x'
2   int value = x;

```

```

3           ^
4 note: became invalid here
5   auto y = std::move(x);
6           ^

```

Step	C++ code	LifetimeCheck actions
1	<code>int x = 42;</code>	AllocaOp : Create <code>x</code> StoreOp : Initialize <code>x</code> <code>pmap[x] = {x}</code> , category: Value
2	<code>int y = std::move(x);</code>	CallOp : Detect <code>std::move</code> call LoadOp : Load from <code>x</code> StoreOp : Store to <code>y</code> Action : Mark <code>x</code> as moved-from <code>pmap[x] = {invalid}</code> , <code>pmap[y] = {y}</code>
3	<code>int z = x;</code>	LoadOp : Attempt to load from <code>x</code> Check : <code>isValueTypeMovedFrom(x)</code> ? YES DIAGNOSTIC : use of moved-from value Note: became invalid at step 2

Figure 5: Step-by-step use-after-move detection example

3 Conclusion

Use-after-move detection needs both sides of ClangIR. AST attributes identify the C++ operation, while the LifetimeCheck state model records which value has become invalid. The next article, [Control Flow, Scope, and Advanced Lifetime Cases](#), extends the same model through memory operations, structured control flow, scopes, coroutines, and return values.

References

- [1] Herb Sutter. Lifetime safety: Preventing common dangling. C++ Standards Committee Paper P1179R1, ISO/IEC JTC1/SC22/WG21, November 2019.