



# Exceptions the Other Way Around

**Sean Parent | Sr. Principal Scientist**  
**Adobe Software Technology Lab**

Artwork by Dan Zucco

Bē

# What is an error?

- Every operation has a set of *preconditions*
- **When preconditions are satisfied, an operation must:**
  - Complete successfully, satisfying *postconditions*
  - Or, return an error with an indication as to why the postconditions could not be satisfied
- There is a logically isomorphic system of thought where postconditions *include* the error state

**Errors are about *Postconditions***



# Preconditions

- A *Precondition* is an assertion that must be true before an operation `sort(first, last, compare)`
  - $[first, last)$  is a valid range (implying  $first \leq last$ )
  - For all  $p$  in the range  $[first, last)$ ,  $p$  is dereferenceable
  - For all  $p$ , let  $v$  equal the set of values  $*p$ ;
    - For all pairs  $(v_a, v_b)$ , *compare* is a predicate establishing a strict-weak-order relation
- The *domain of an operation* is the set of values satisfying all preconditions

# Postconditions

- A *Postcondition* is an assertion that must be true just after an operation
  - Unless there is an error
- The postcondition of `sort()` is that all the elements in the range `[first, last)` are in non-decreasing order as defined by `compare`

# Class Invariants

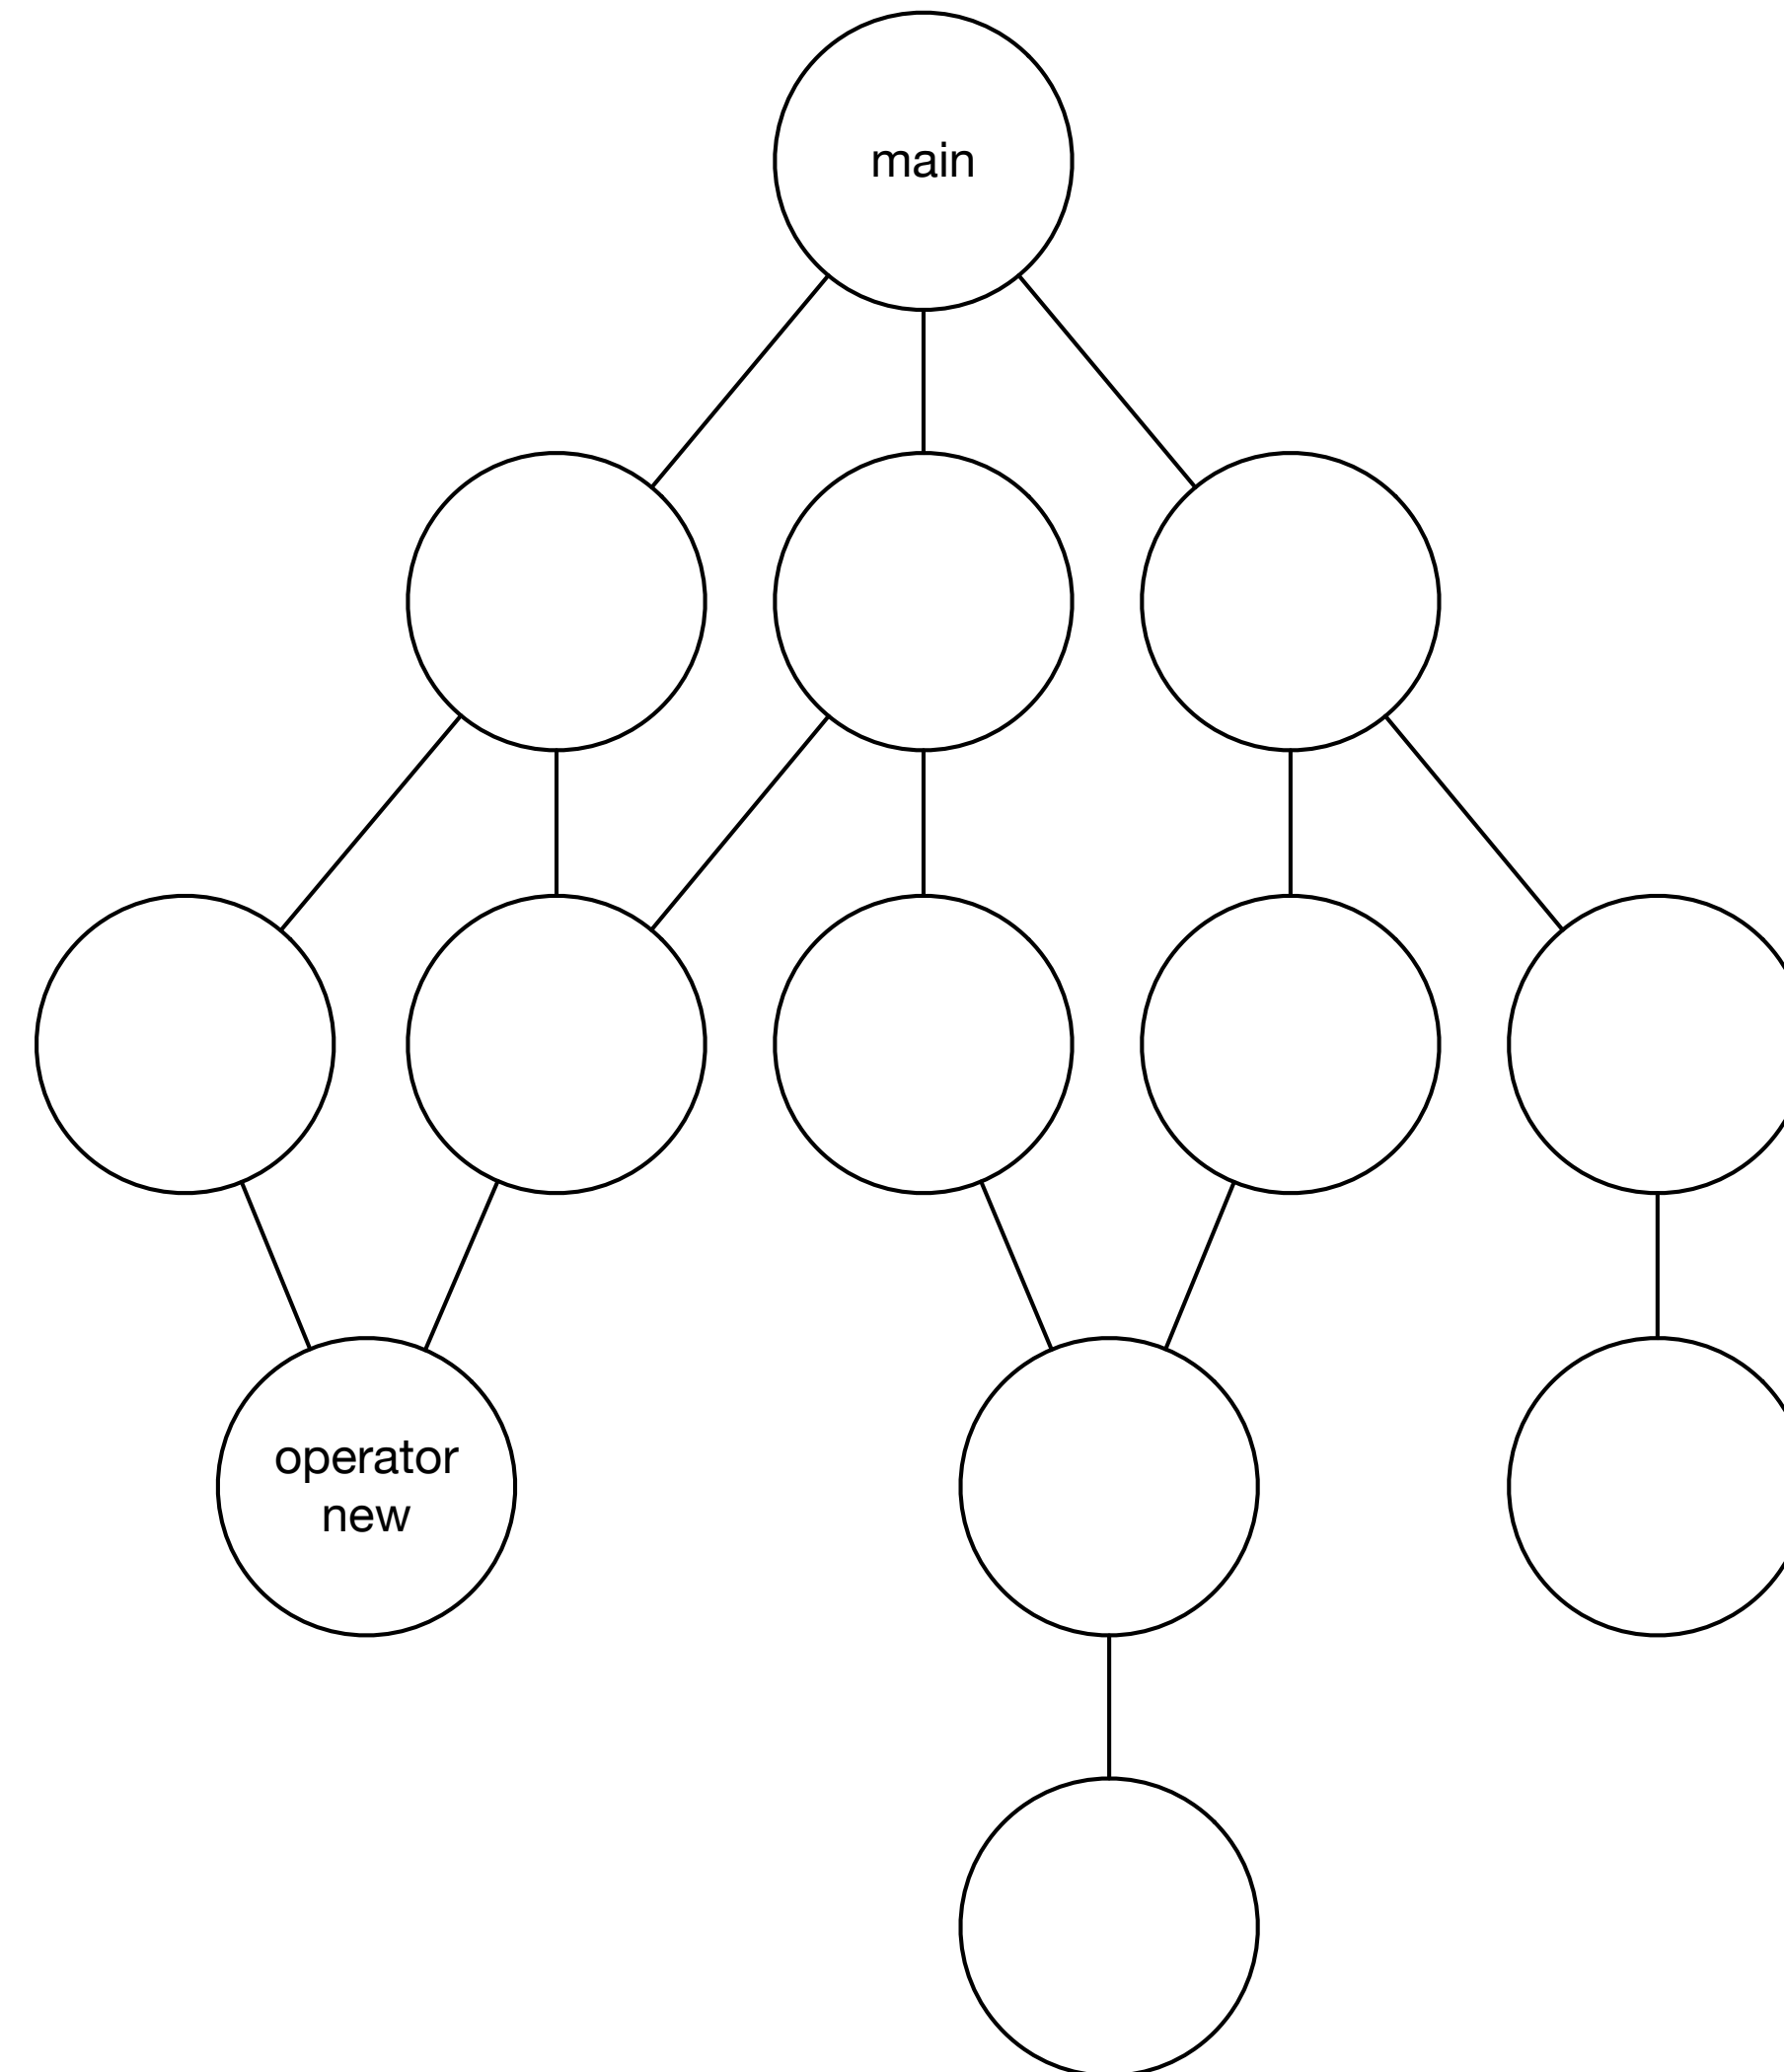
- *Class invariants* are postconditions that hold for all operations on a type
- As such, they can be counted on to hold as a precondition for all operations and don't need to be explicitly stated.



# When to report an error

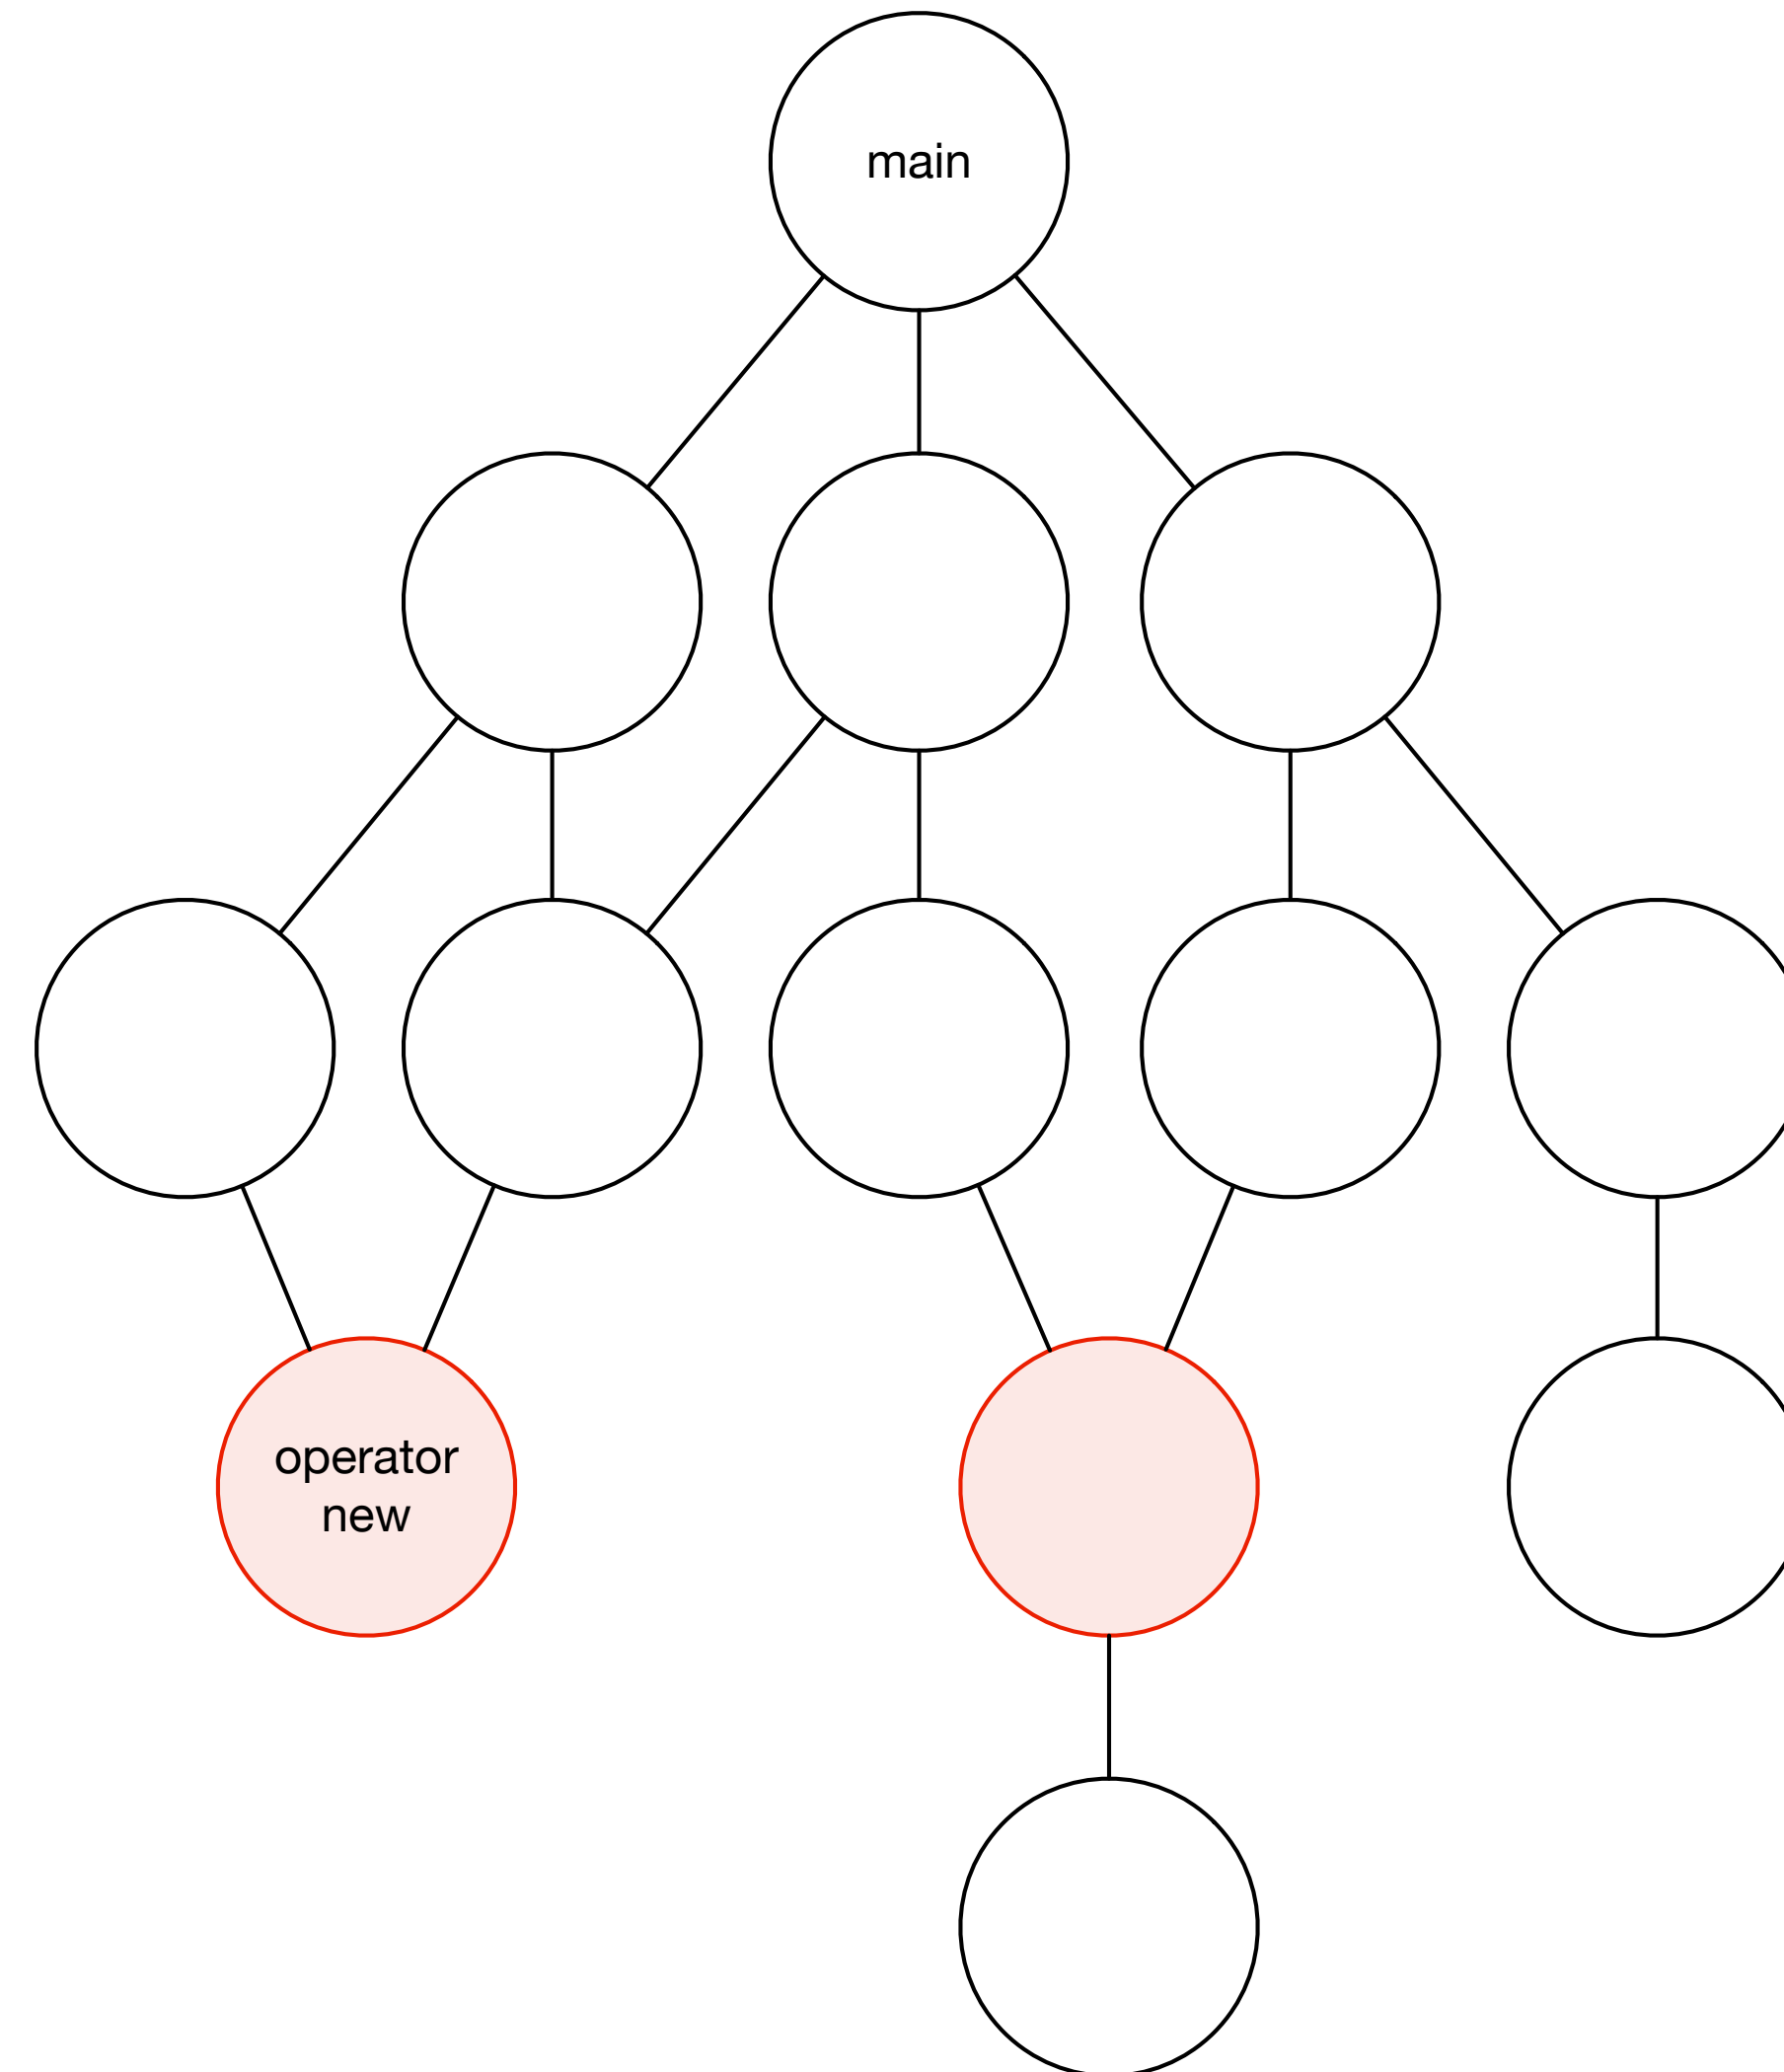
- An error is appropriate when a **postcondition** cannot be satisfied
  - An error is a *recoverable* event
  - Programming errors are not recoverable because you can't tell where they came from
- Resource exhaustion (i.e. out-of-memory)
- I/O failure
- Validating external data
- Cancellation
- Implementation and representation limits

# Errors tend to happen at a low level





# Errors tend to happen at a low level




# Key Points

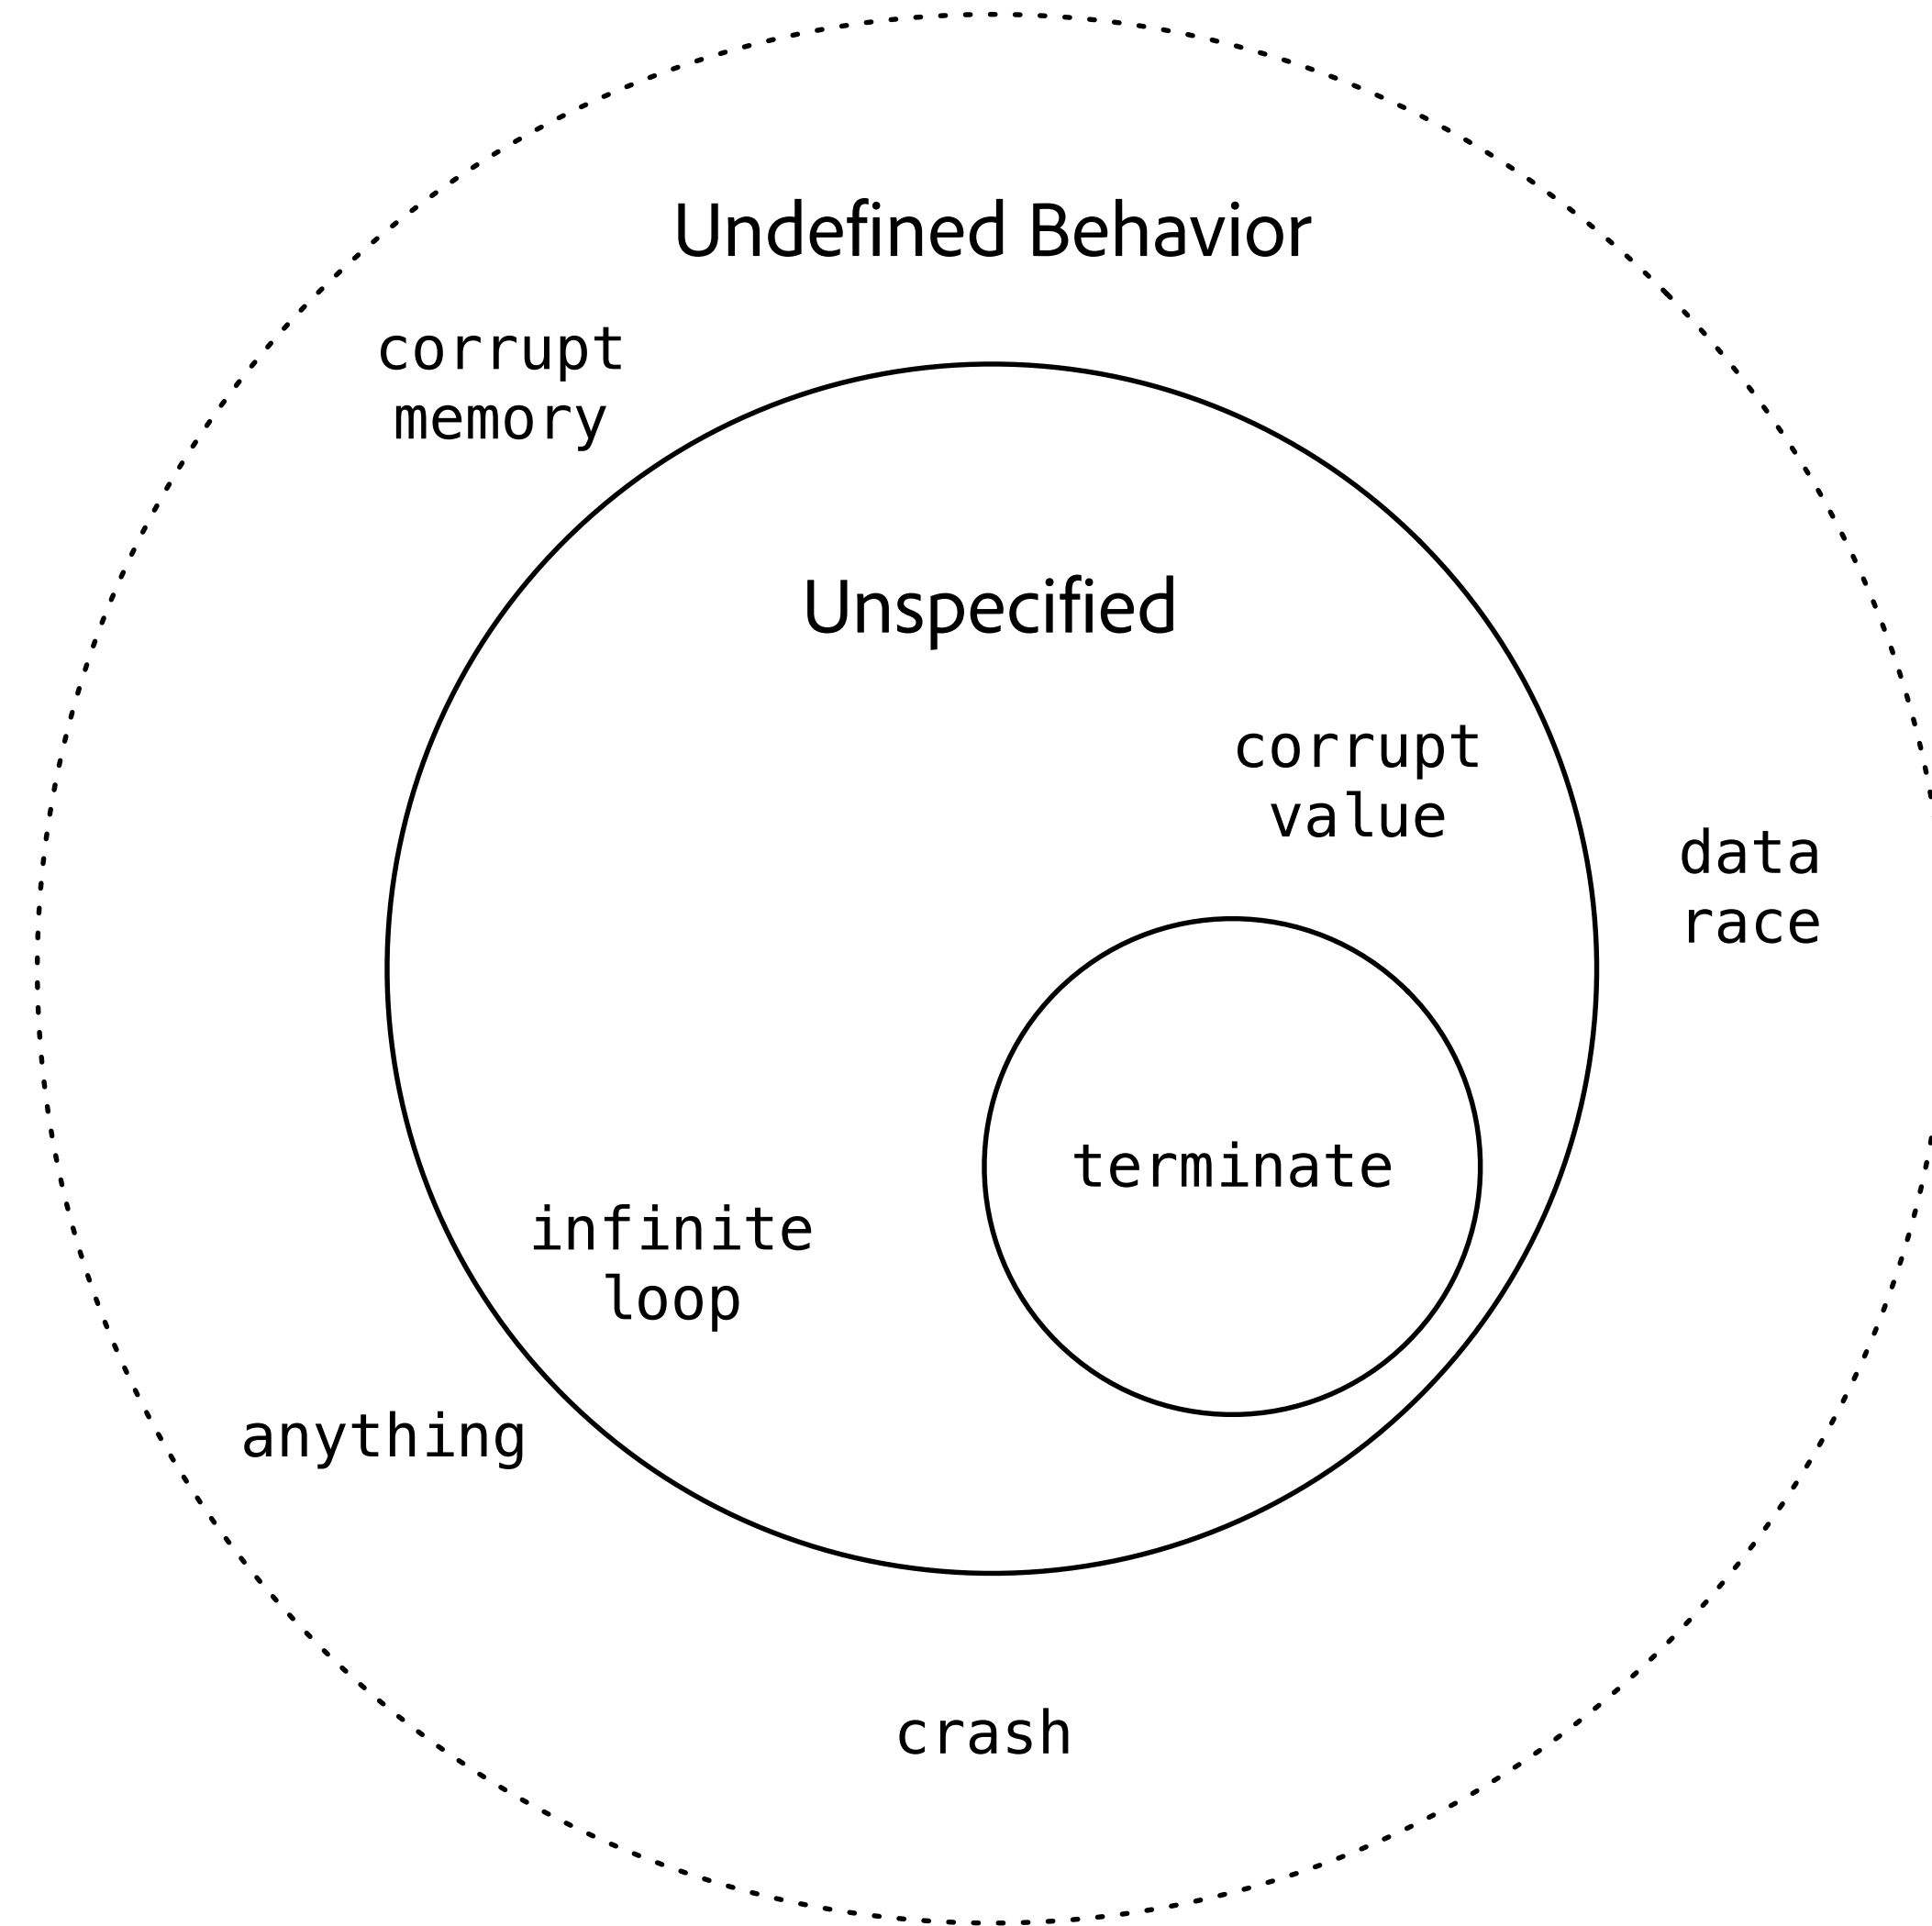
- Errors tend to occur at a low-level
- A significant amount of code may be in the path from an error to the point where it can be handled

# Preconditions

- **When preconditions are not satisfied:**
  - An operation may lead to *undefined behavior*
  - The result may be *unspecified* and may violate *class invariants*
  - It may lead to program *termination*



# Possible Effects of Precondition Violation



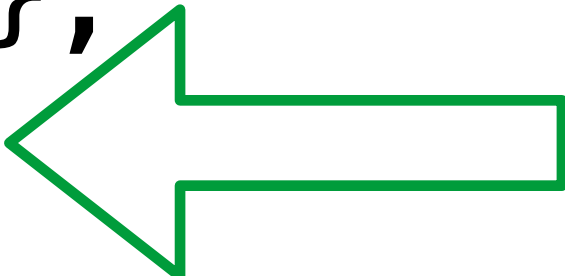
# Safety

- An operation is *safe* if it cannot lead to undefined behavior
  - Directly or indirectly
  - Even if the operation preconditions are violated
- An *unsafe* operation may lead to undefined behavior if preconditions are violated
  - Either directly or during subsequent operations, safe or not
- We refer to an operation that terminates on a precondition violation or has no preconditions, as *strongly safe*

# Safety

- Safety is about incorrect code and the scope of damage it may cause
- Errors are about correct code and recoverable situations
  
- Safety is a transitive property
- Correctness is not transitive
- Strong safety is not transitive

# What is *valid*?


```
int* p{nullptr};  
{  
    int x{0};  
    p = &x;   
}
```

- `p` is *valid* and dereferencable



# What is *valid*?


```
int* p{nullptr};  
{  
    int x{0};  
    p = &x;  
    ++p;  
}
```



- `p` is *valid* but not dereferencable
- `*p` is undefined behavior

# What is *valid*?


```
int* p{nullptr};  
{  
    int x{0};  
    p = &x;  
    ++p;  
}  
// p
```



- `p` is *invalid*, `p` may be assigned-to, or destructed
- `*p` is undefined behavior
- all other operations, including copy and comparisons, are implementation-defined and may trap

## What is *valid*?

```
int* p{nullptr};  
{  
    int x{0};  
    p = &x;  
    ++p;  
}  
p = nullptr;
```



- `p` is valid but not dereferencable

# Myths

“It's almost as though exceptions are viewed as a mysterious attack on otherwise correct code, from which we must protect ourselves. Needless to say, this doesn't lead to a healthy relationship with error handling!”

# Guarantees

- *Basic exception guarantee*
  - *Invariants* hold
  - The values of objects being modified are otherwise unspecified
- *Strong exception guarantee*
  - The state is restored to the state prior to the failing operation
  - *Invariants* hold by extension
- *No exception guarantee*

# Fix mutating state when propagating an exception

- The exception guarantees are about what happens to mutable state
- Operations that do not mutate state can propagate errors directly

# Class Invariants & Error Propagation

- The guarantees say we need to consider:
  - Operations where invariants are temporarily broken
- Only consider member functions that modify the object



# vector example

- Consider
  - operator=
  - assign
  - reserve
  - shrink\_to\_fit
  - clear
  - insert
  - erase
  - push\_back
  - emplace\_back
  - pop\_back
  - resize
- Ignore
  - constructors
  - get\_allocator
  - at, operator[]
  - front
  - back
  - data
  - begin, end
  - empty
  - size
  - max\_size
  - capacity
  - comparison operators
  - destructor

# vector example

- We also don't have to worry about non-*basis operation*

```
void resize(size_type sz, const T& c) {
    if (sz < size()) {
        erase(begin() + sz, end());
    } else {
        T t{c};
        reserve(sz);
        for (size_type f = 0, l = size() - sz; f != l; ++f) {
            push_back(t);
        }
    }
}
```

# Fix invariants when propagating an exception

- Invariants should only be broken inside member functions with private access
- Keep your *basis* small to narrow the impact of exceptions

# Example

```
template <class T, class U>
class zip_vector {
    // invariant: _v0.size() == _v1.size()
    vector<T>    _v0;
    vector<U>    _v1;
public:
    void push_back(T&& x, U&& y) {
        _v0.push_back(move(x));
        _v1.push_back(move(y)); // what if this throws?
    }
};
```

# Example

```
template <class T, class U>
class zip_vector {
    // invariant: _v0.size() == _v1.size()
    vector<T>    _v0;
    vector<U>    _v1;
public:
    void push_back(T&& x, U&& y) {
        _v0.push_back(move(x));
        try {
            _v1.push_back(move(y));
        } catch(...) {
            _v0.pop_back();
            throw;
        }
    }
};
```

# Postconditions

- A *Postcondition* is an assertion that must be true just after an operation
  - Unless there is an error
    - Then invariants are satisfied but the state is otherwise unspecified
- The basic guarantee falls out of all the documentation of your system

# Safety and Exception Guarantees

- The basic and strong exception guarantees are statements about correctness
  - They are not transitive properties
  - “It is ‘safe’ in the sense that it is not allowed to crash, but its output may be unpredictable.”
    - David Abrahams regarding the basic guarantee



# The Other Way Around

- Top-down vs bottom-up design
- In Generic Programming, *lifting* is finding commonalities bottom-up across component
- The *other way around* is to find common *requirements* of callers and common models

# Requirements

- Generalization of preconditions to also include
  - Operations & semantics
  - Dependent types
- In non-generic code, these requirements exist but are implied by the type
- Concepts allow us to specify type requirements
  - But are still missing value preconditions

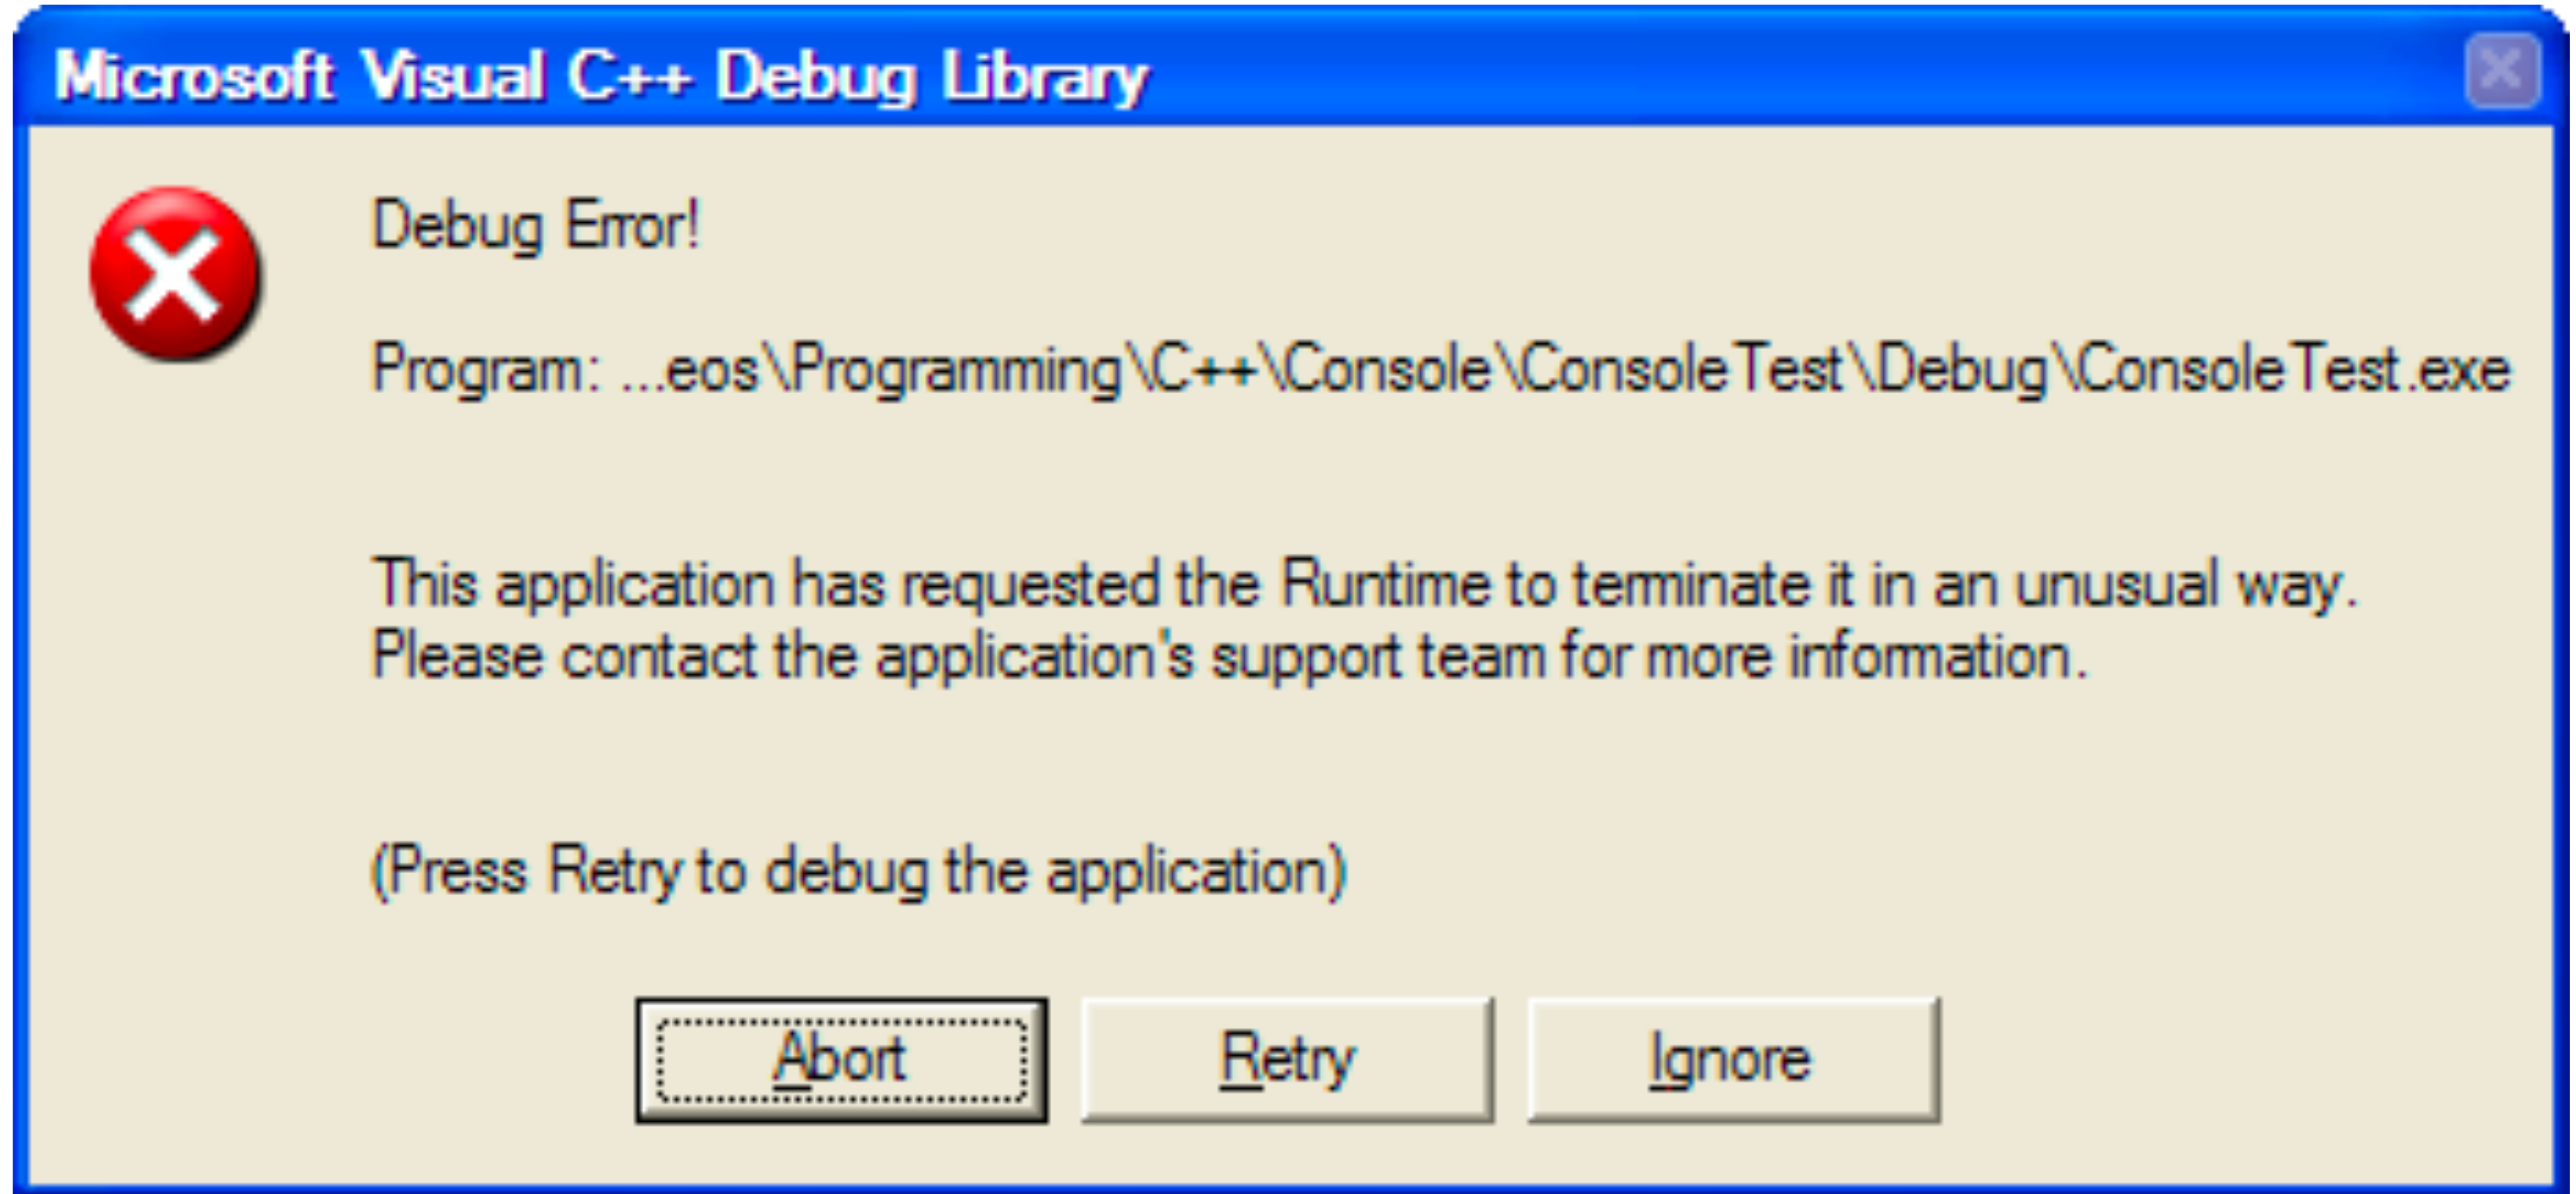
# Guarantees

- Generalization of postconditions
  - May be conditionalized

# Example

- `sort` **requires** that `value_type(iterator)` is movable
- `basic_string<T>` **guarantees** it is movable
- therefore an `array<string>` is sortable

# Stopping Error Propagation



# Stopping Error Propagation

- The operation must satisfy its postconditions
  - Or terminate

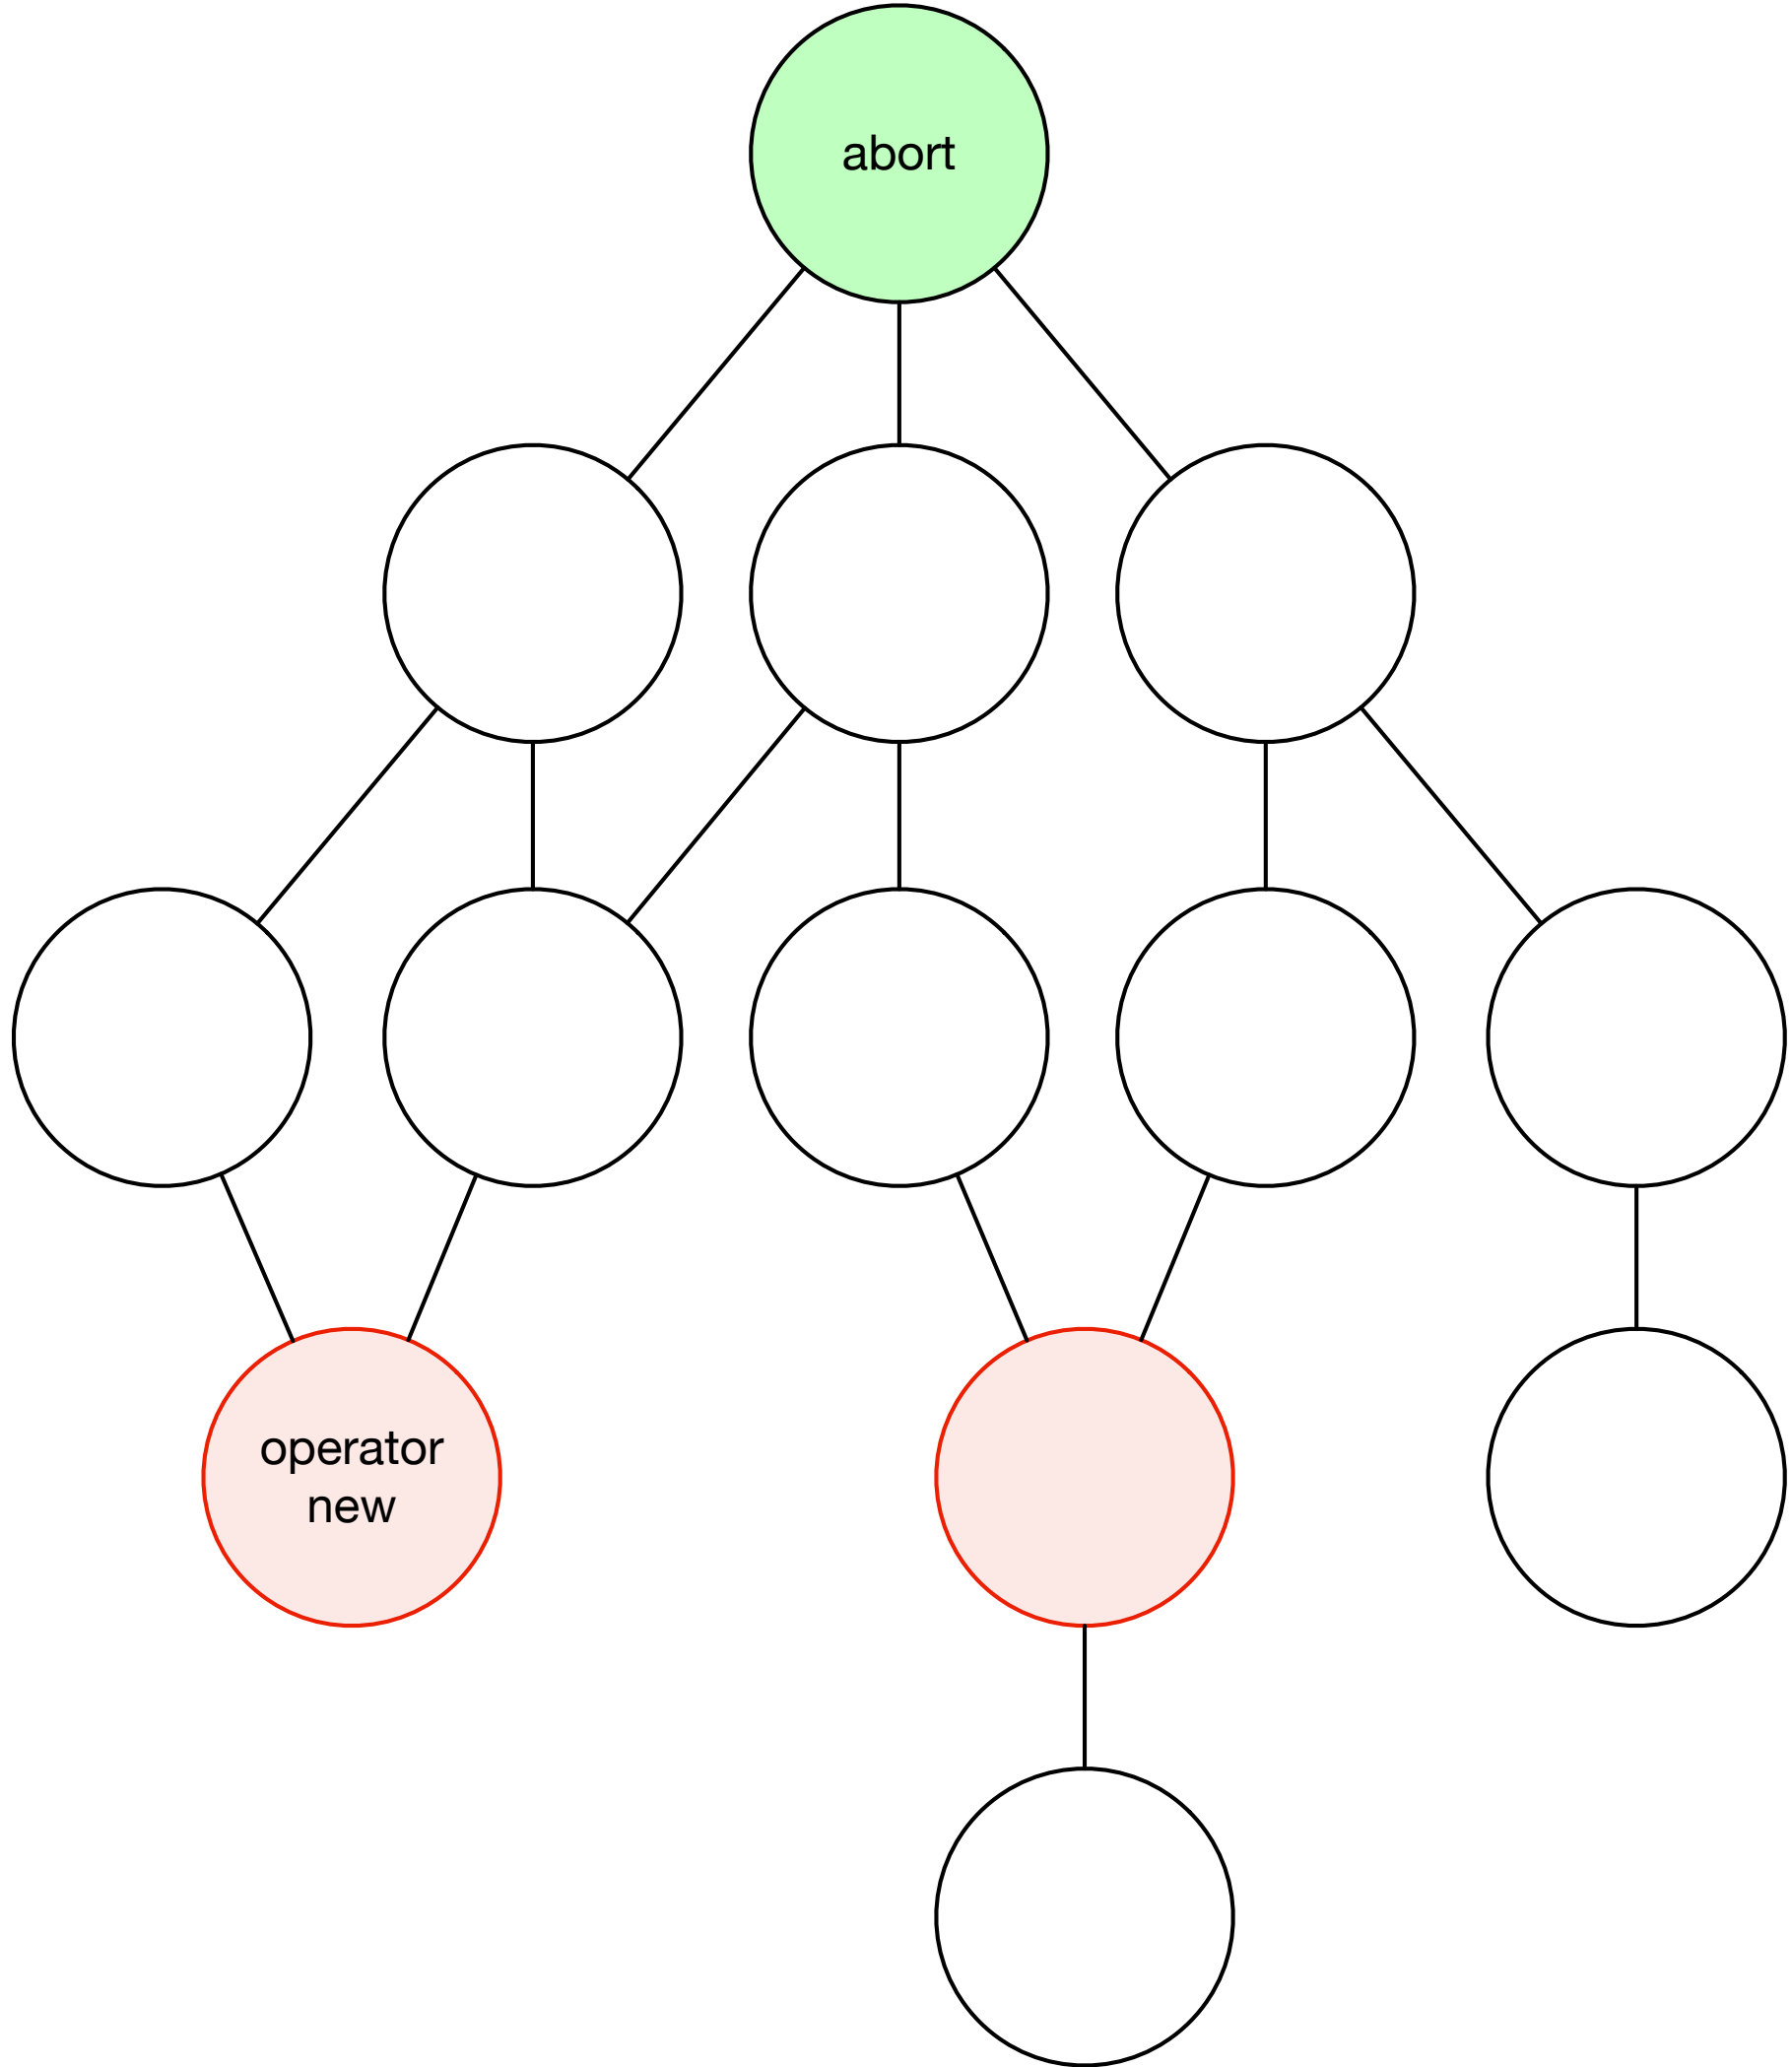
# Abort



# Report & Terminate

- Program cannot satisfy postconditions
- Requirements
  - None
- Appropriate for failure during initialization

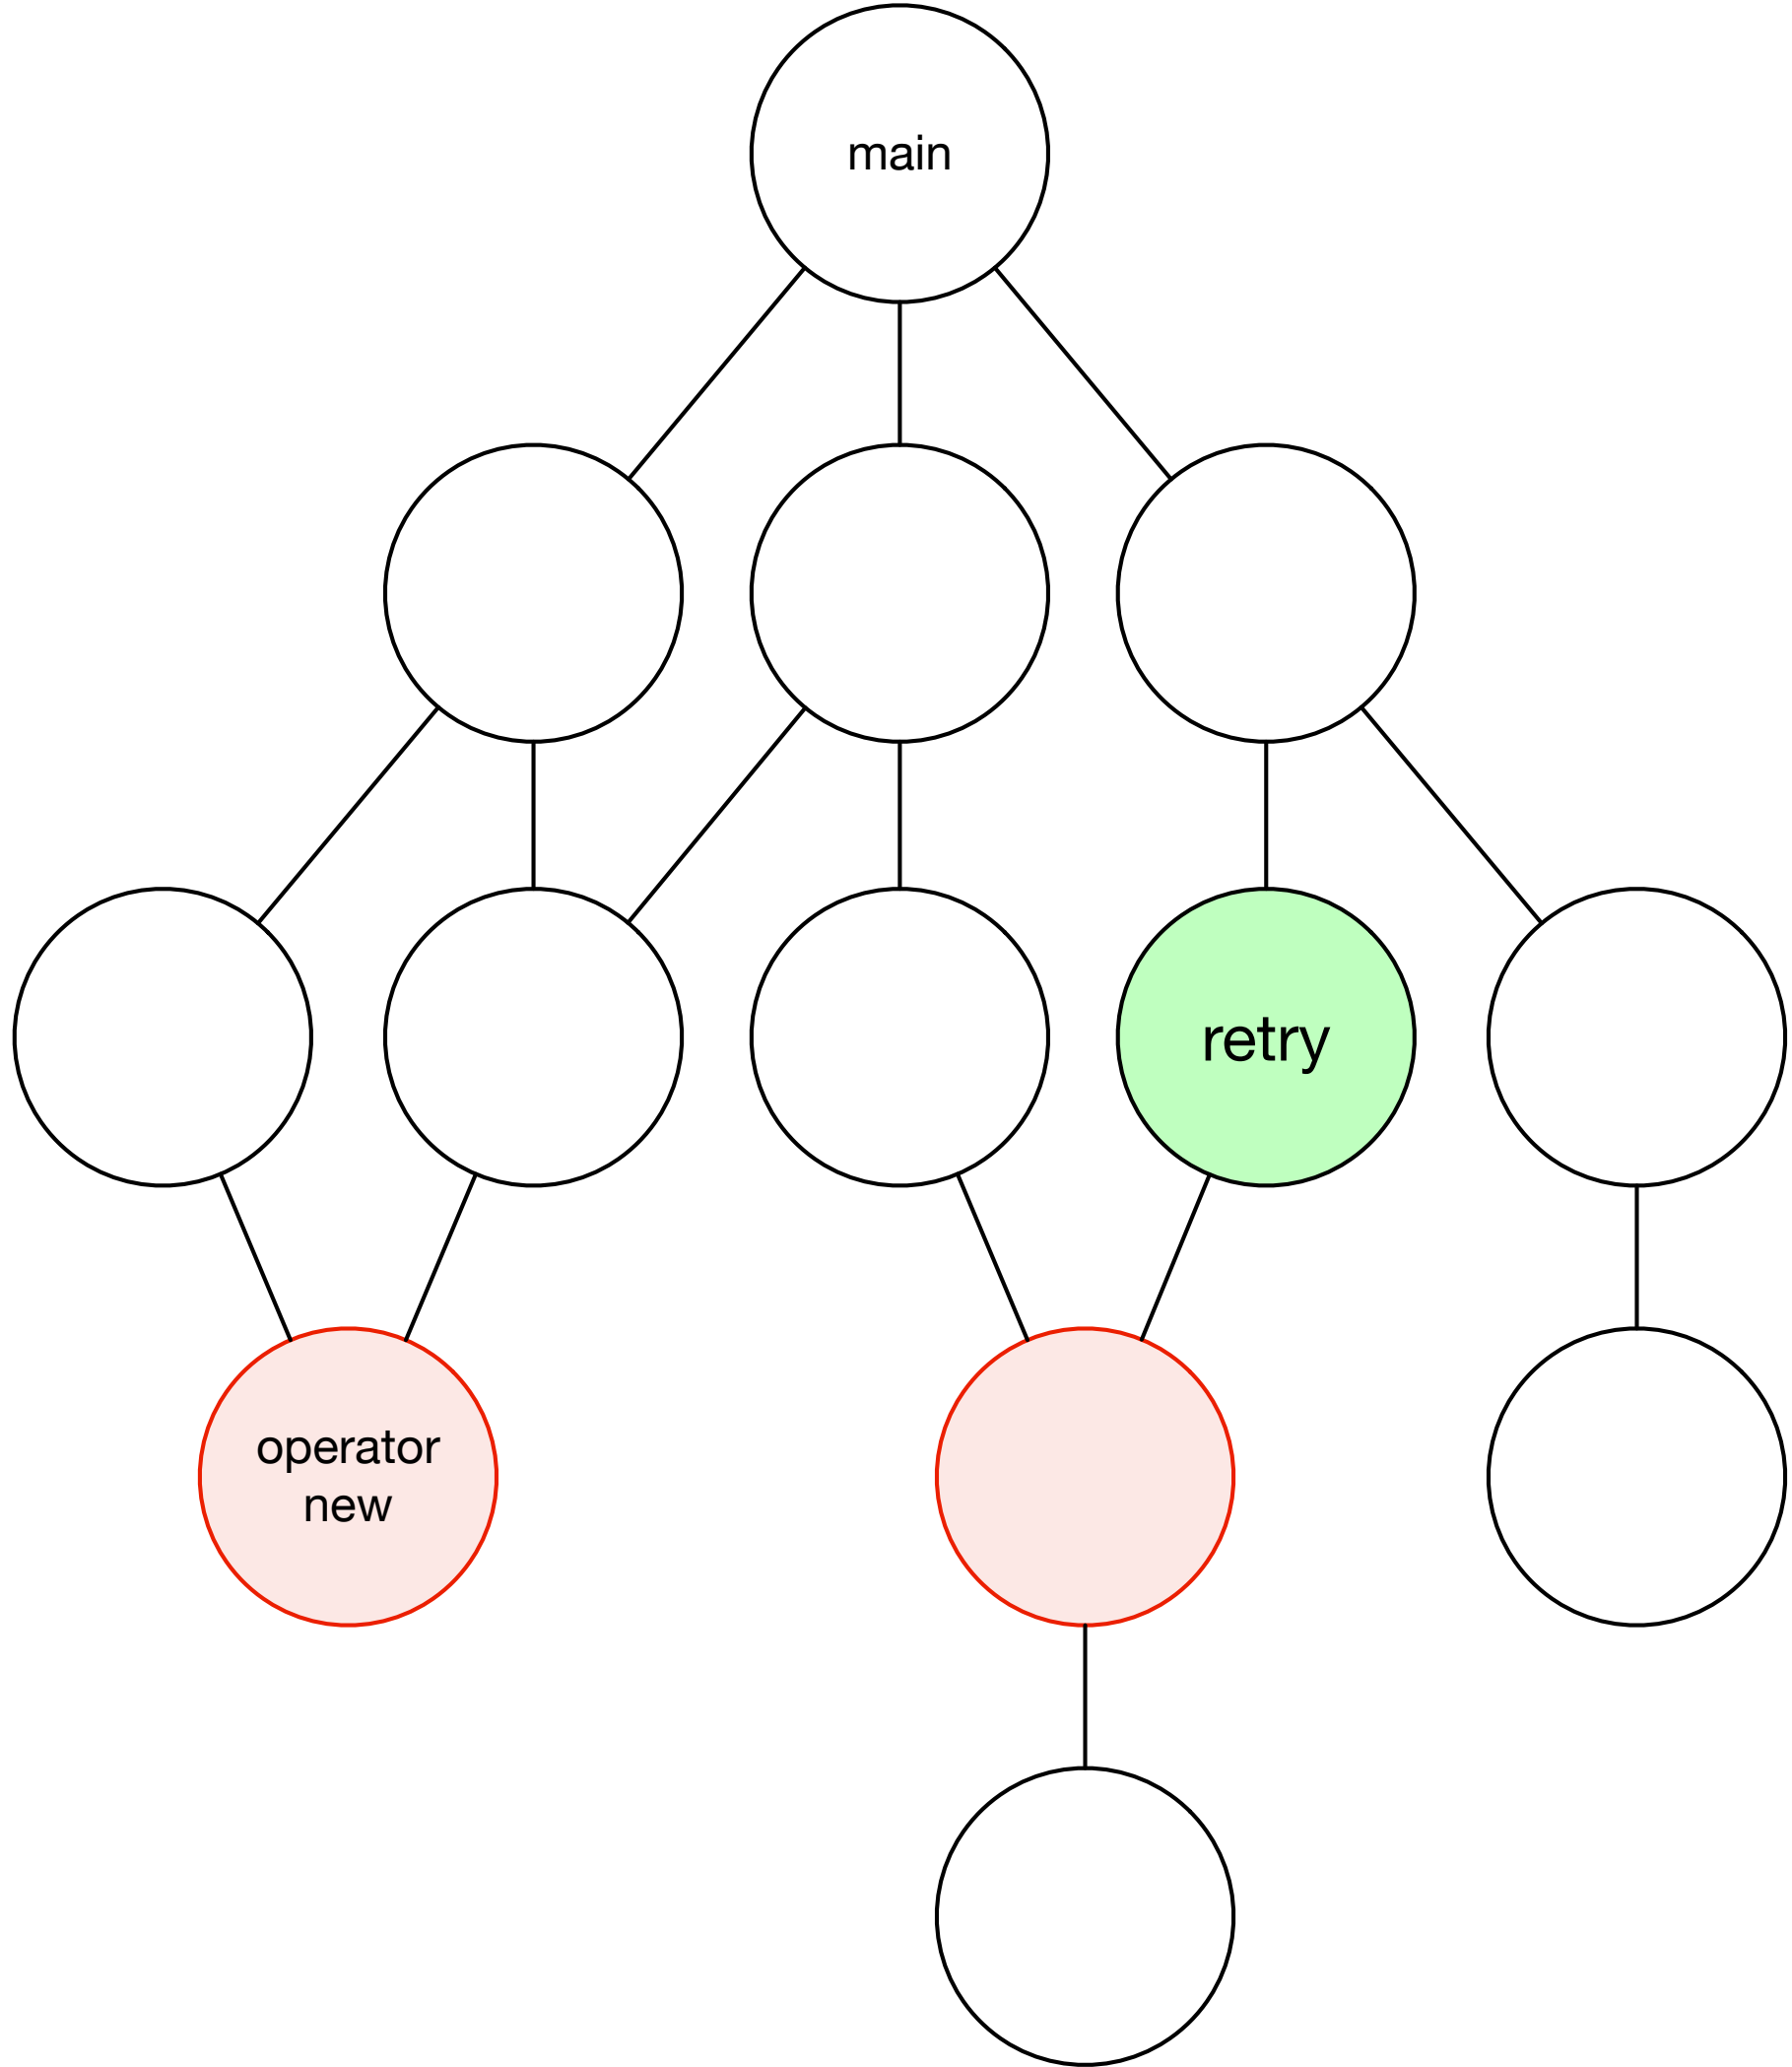
# Abort is Likely to Happen Near Main



# Retry

- Retry the operation using the same, or a different, approach
  - Successful retry will satisfy postconditions
  - Or, resume error propagation
- Requirements
  - Discard incomplete work
- Appropriate for some I/O, memory cache
- Best of done close to point of failure

# Retry is Best Handled Near the point of Failure

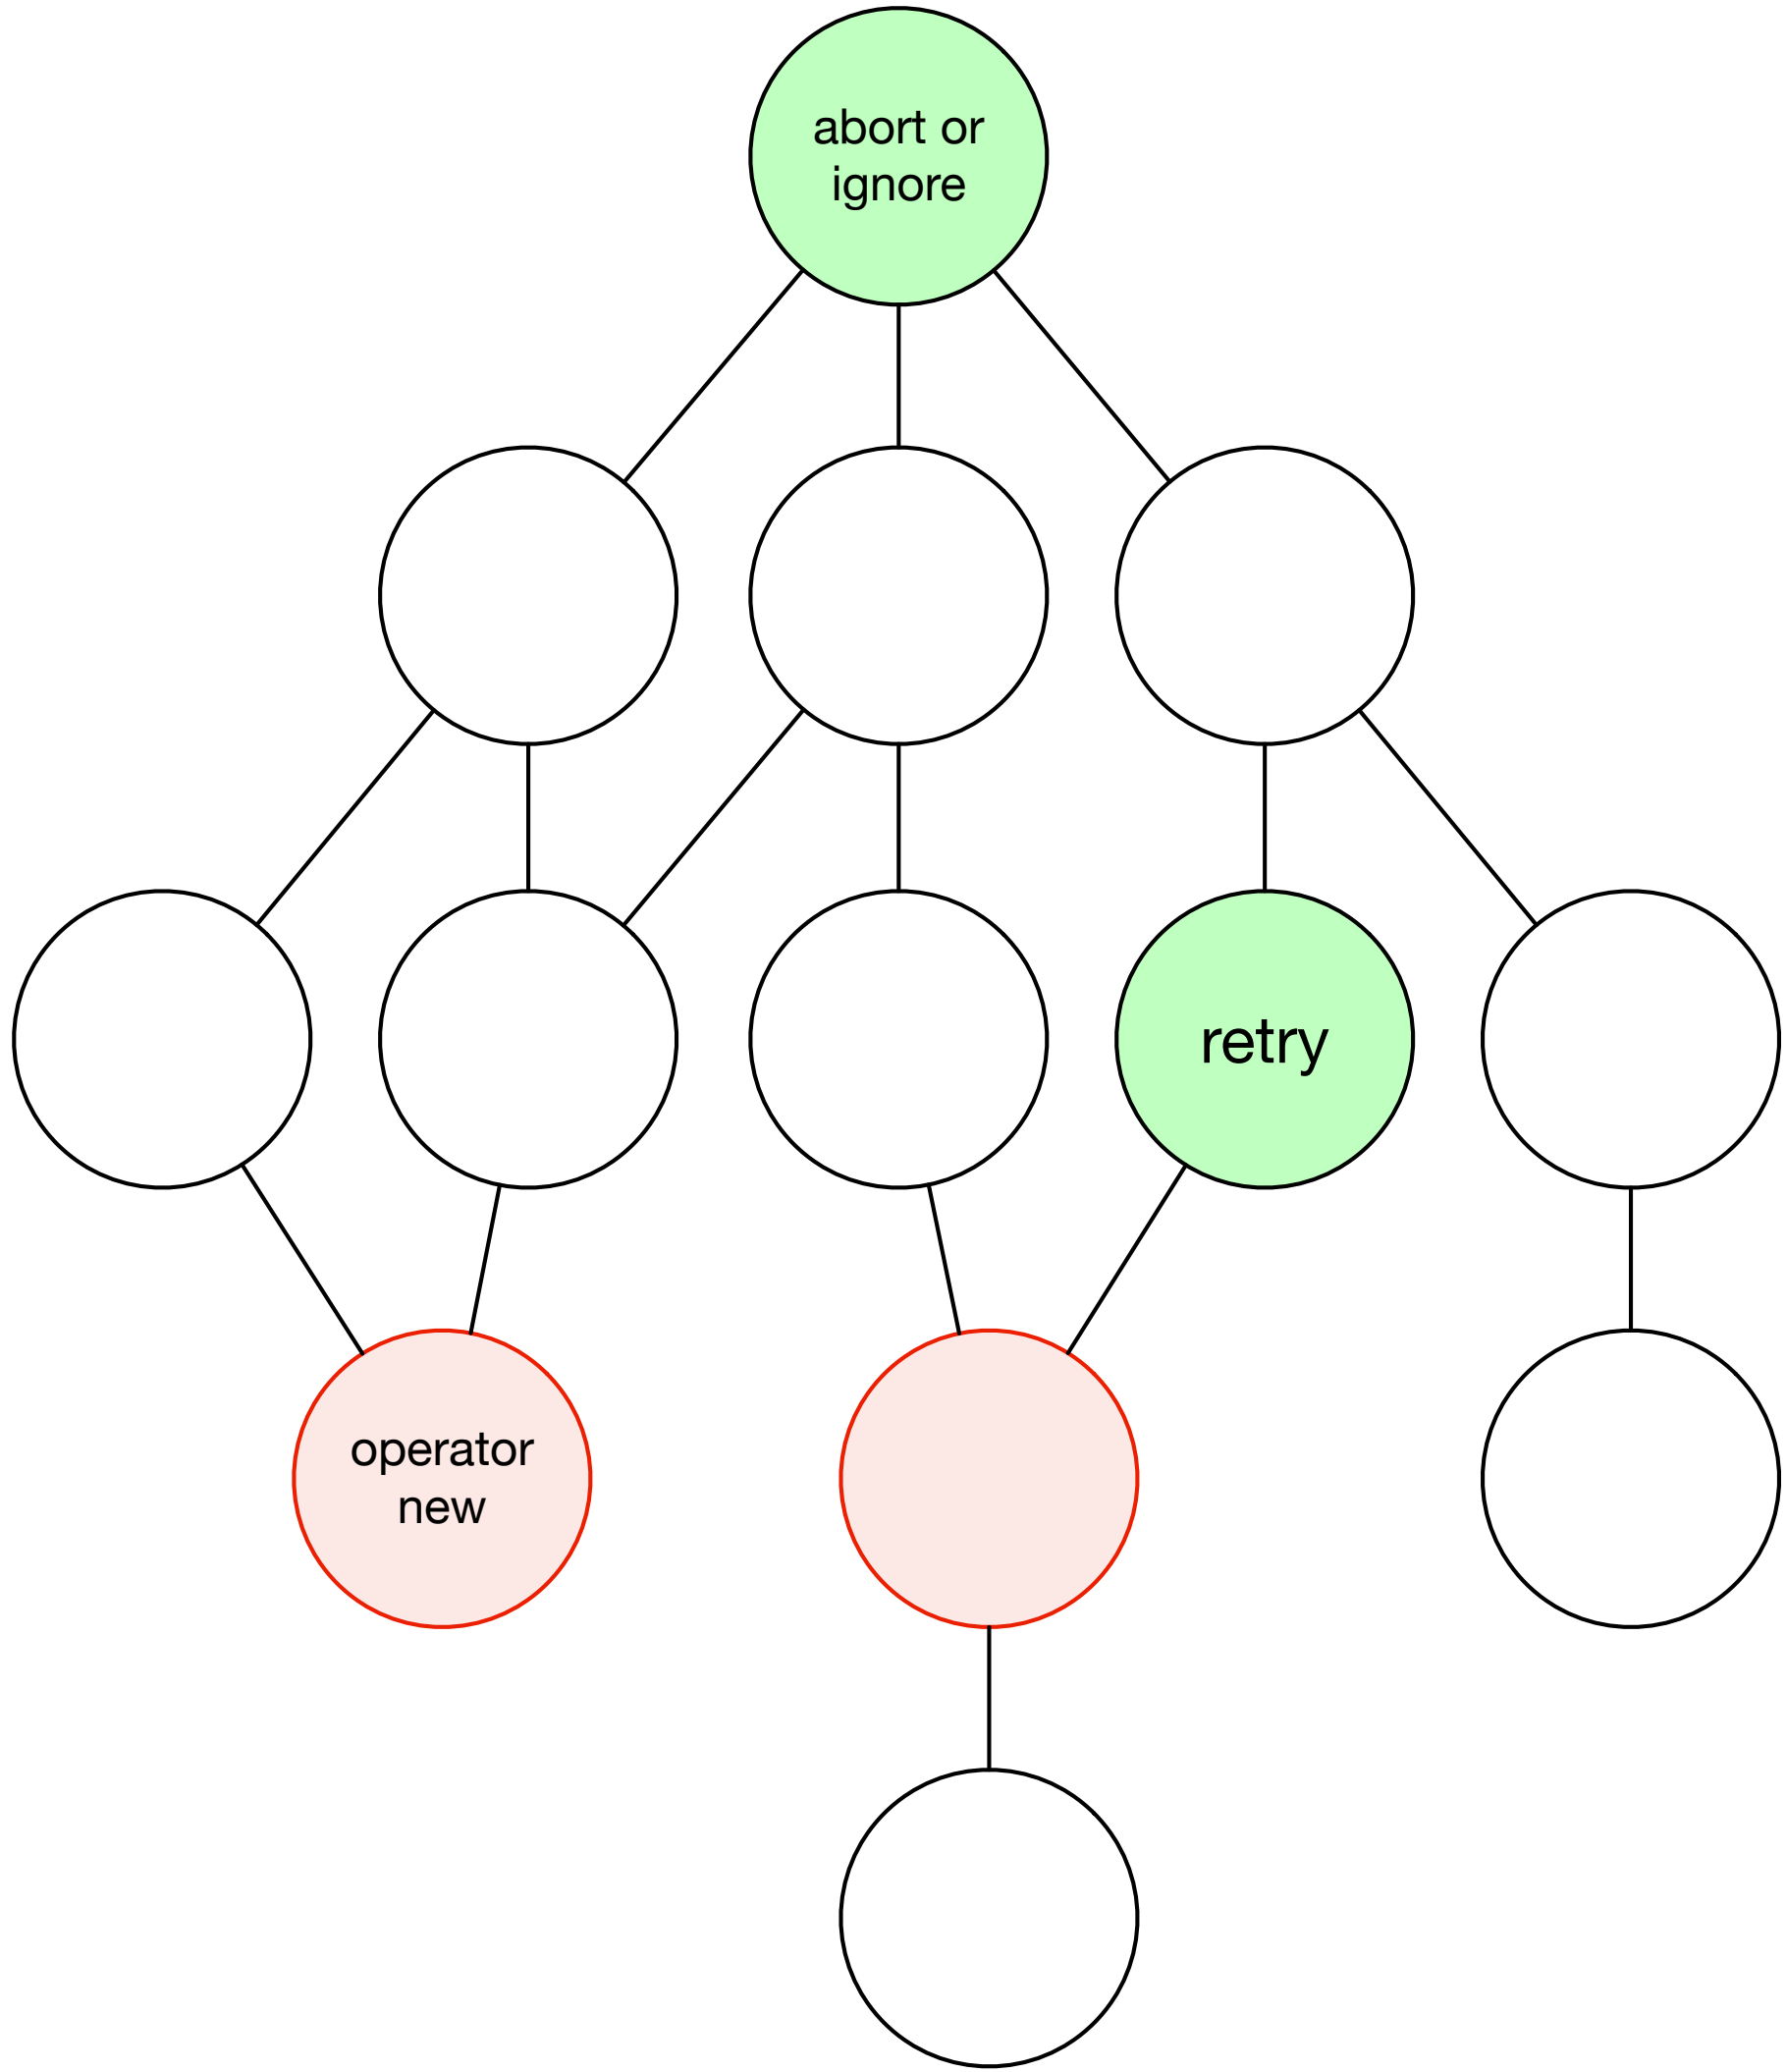


# Ignore

# Report & Continue

- Requirements
  - Discard incomplete work
- Often implemented as *transactional* operations
  - Strong exception guarantee
- Appropriate for sequenced or optional operations

# Retry is Best Handled Near the Point of Failure



# Discard Incomplete Work

- “The options for recovery, in this case, are limited: either destruction or resetting the component to some known state before further use.”
  - Dave Abrahams regarding basic exception guarantee



# Requirements of Exceptions

- What do we mean by “discard partial work”?
- An object which is being mutated when an exception is thrown must leave the object in a state
  - within the domain of destruction
  - within the domain of the left-hand side of assignment
- The object should also be in the domain of any operation which does not read the object
  - i.e. `clear()`

# Requirements of Exceptions

- The exception guarantees are not requirements
- The implied requirement is “satisfies class invariants”
  - But that doesn’t actually tell you anything unless there is a requirement that some operations have no preconditions

# Invariants?

- We'll call the class invariants in the absence of exceptions, the *desired invariants*
- The weaker invariant of only destructible and assignable-to is *invalid*
  
- We can, trivially, weaken class invariants to include this state
  - `invalid() || desired_invariants()`
  
- And add the precondition for all the other operations
  - `valid() && ...`

# Why invalid?

- Destructible and assignable-to are transitive
  - To destruct an object, destruct all the parts
  - To assign an object, assign all the parts
- If we define the class invariants and preconditions of its operations to include invalid
  - Without doing anything else in code to account for exceptions
  - A class using “= default” for destruction and assignment can be *exception correct*

# Other Cases Simplified

- Ensuring destructible and assignable-to in cases where the default implementations are insufficient is often simpler than reasoning about the desired invariants
- Parts can be packaged to allow default definitions
  - i.e. `unique_ptr<T>` vs `T*`

# Relationship to Move

- Any function operation of the form  $x = f(x)$  can be expressed as a mutating operation  $a(x)$ .
- It can also be expressed as a consuming operation  $x = g(\text{move}(x))$
- If  $g()$  throws an exception,  $x$  is in a moved-from state
- The required postcondition after an error are identical to those of a moved-from object
  - See P2345

# Class Invariants & Postconditions

- A *class invariant* is a postcondition
- A *postcondition* is an assertion that must be true just after an operation
  - Unless there is an error
    - Then the result may be invalid

# Example

```
template <class T, class U>
class zip_vector {
    // invariant: invalid() || (_v0.size() == _v1.size())
    vector<T>    _v0;
    vector<U>    _v1;
public:
    // precondition: valid()
    void push_back(T&& x, U&& y) {
        _v0.push_back(move(x));
        _v1.push_back(move(y));
    }
};
```



# Example

```
template <class T, class U>
class zip_vector {
    // invariant: _v0.size() == _v1.size()
    // may be invalid by exception
    vector<T>    _v0;
    vector<U>    _v1;
public:
    void push_back(T&& x, U&& y) {
        _v0.push_back(move(x));
        _v1.push_back(move(y));
    }
};
```

**Correct Exception Handling has  
Zero Impact on Most Code**







Adobe

Bē

Artwork by Dan Zucco



# Disadvantage

- Stronger preconditions may result in unsafe operations

# Addendum 1: What about shared state?

- Correct mutation requires exclusivity
- Everything within the program is part of the whole program
- You must understand the implicit structure as a whole
  - The scope of which is somewhere between the calling scope and the program as a whole
  - You cannot reason about it locally
- By catching close to `main()` you may be *under* the implicit structure
  - Or be able to identify the root(s)

# Addendum 2: The Problems with C++ Exceptions

- Exceptions are the default
- Operator new throws and is replaceable
- Reliance on RTTI and inheritance
- No guarantee of `noexcept` move for standard components
- ABI leakage

## About the artist

### Dan Zucco

London-based 3D art and motion director Dan Zucco creates repeating 2D patterns and brings them to life as 3D animated loops. Inspired by architecture, music, modern art, and generative design, he often starts in Adobe Illustrator and builds his animations using Adobe After Effects and Cinema 4D. Zucco's objective for this piece was to create a geometric design that felt like it could have an infinite number of arrangements.

Made with

 Adobe Illustrator

 Adobe After Effects







Adobe

Bē

Artwork by Dan Zucco