

No littering!



Bjarne Stroustrup

Morgan Stanley, Columbia University

www.stroustrup.com


Executive summary

- We now offer complete type- and resource-safety
 - No memory corruption
 - No resource leaks
 - No garbage collector (because there is no garbage to collect)
 - No runtime overheads (Except where you need range checks)
 - No new limits on expressibility
 - ISO C++ (no language extensions required)
 - Simpler code
- Support
 - C++ Core Guidelines: <https://github.com/isocpp/CppCoreGuidelines>
 - GSL: <https://github.com/microsoft/gsl>
 - Static analysis/enforcement: In Microsoft Visual Studio
- We want “C++ on steroids”
 - Not some neutered subset





I ♥ C++

Caveat: Not yet deployed at scale ☹️



C++ use

- About 4.5M C++ developers
- 2014-2015: increase of 200,000 developers
- www.stroustrup.com/applications.html




The big question

- “What is good modern C++?”
 - *Many* people want to write “Modern C++”
- What would you like your code to look like in 5 years time?
 - “Just like what I write today” is a poor answer
- Guidelines project
 - <https://github.com/isocpp/CppCoreGuidelines>
 - Produce a *useful* answer
 - Implies tool support and enforcement
 - Enable *many* people to use that answer
 - For most programmers, not just language experts





Overview

- Pointer problems
 - Memory corruption
 - Resource leaks
 - Expensive run-time support
 - Complicated code
- The solution
 - Eliminate dangling pointers
 - Eliminate resource leaks
 - Library support for range checking and **nullptr** checking
 - And then deal with casts and unions




I like pointers!

- Pointers are what the hardware offers
 - Machine addresses
 - For good reasons
 - They are simple
 - They are general
 - They are fast
 - They are compact
- C's memory model has served us really well
 - Sequences of objects
- But pointers are not “respectable”
 - Dangerous, low-level, not mathematical, ...
 - There is a huge ABP crowd




Lifetime can be messy



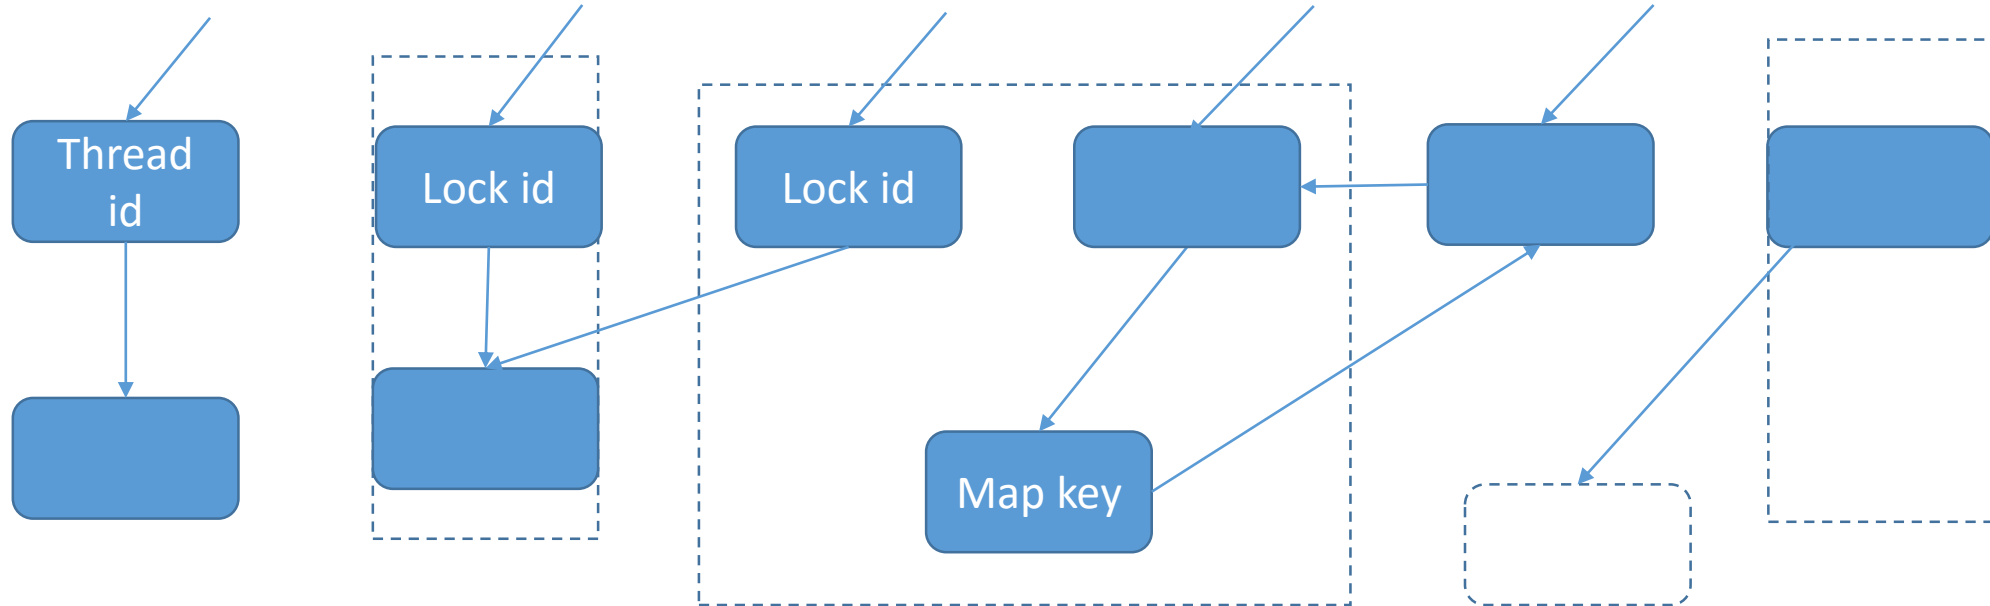
- An object can have
 - One reference
 - Multiple references
 - Circular references
 - No references (leaked)
 - Reference after deletion (dangling pointer)

Ownership can be messy




- An object can be
 - on stack (automatically freed)
 - on free store (must be freed)
 - in static store (must never be freed)
 - in another object

Resource management can be messy



- Objects are not just memory
- Sometimes, significant cleanup is needed
 - File handles
 - Thread handles
 - Locks
 - ...

Access can be messy



- Pointers can
 - point outside an object (range error)
 - be a **nullptr** (useful, but don't dereference)
 - be uninitialized (bad idea)

Eliminate all leaks and all memory corruption

- Every object is constructed before use
 - Once only
- Every fully constructed object is destroyed
 - Once only
 - Every object allocated by **new** must be **deleted**
 - No scoped object must be **deleted** (it is implicitly destroyed)
- No access through a pointer that does not point to an object
 - Read or write
 - Off the end of an object (out of range)
 - To **deleted** object
 - To “random” place in memory (e.g., uninitialized pointer)
 - Through **nullptr** (originally: “there is no object at address zero”)
 - That has gone out of scope

Current (Partial) Solutions

- Ban or seriously restrict pointers
 - Add indirections everywhere
 - Add checking everywhere
- Manual memory management
 - Combined with manual non-memory resource management
- Garbage collectors
 - Plus manual non-memory resource management
- Static analysis
 - To supplement manual memory management
- “Smart” pointers
 - Starting with counted pointers
- Functional Programming
 - Eliminate pointers




Current (Partial) Solutions

- These are old problems and old solutions
 - 40+ years
- Manual resource management doesn't scale
- Smart pointers add complexity and cost
- Garbage collection is at best a partial solution
 - Doesn't handle non-memory solutions ("finalizers are evil")
 - Is expensive at run time
 - Is non-local (systems are often distributed)
 - Introduces non-predictability
- Static analysis doesn't scale
 - Gives false positives (warning of a construct that does not lead to an error)
 - Doesn't handle dynamic linking and other dynamic phenomena
 - Is expensive at compile time



Constraints on the solution

- I want it ***now***
 - I don't want to invent a new language
 - I don't want to wait for a new standard
- I want it guaranteed
 - "Be careful" isn't good enough
- Don't sacrifice
 - Generality
 - Performance
 - Simplicity
 - Portability



A solution

- Be precise about ownership
 - Don't litter
 - Offer static guarantee
- Eliminate dangling pointers
 - Static guarantee
- Make general resource management implicit
 - Hide every explicit delete/destroy/close/release
 - “lost of explicit annotations” doesn't scale
 - becomes a source of bugs
- Test for **nullptr** and range
 - Minimize run-time checking
 - Use checked library types
- Avoid other problems with pointers
 - Avoid cast and un-tagged unions



No resource leaks

- We know how
 - Root every object in a scope
 - `vector<T>`
 - `string`
 - `ifstream`
 - `unique_ptr<T>`
 - `shared_ptr<T>`
 - `lock_guard<T>`
 - RAII
 - “No naked **new**”
 - “No naked **delete**”
 - Constructor/destructor
 - “since 1979, and still the best”



Dangling pointers – the worst problem

- One nasty variant of the problem

```
void f(X* p)
{
    // ...
    delete p;           // looks innocent enough
}

void g()
{
    X* q = new X;        // looks innocent enough
    f(q);
    // ... do a lot of work here ...
    q->use();             // Ouch! Read/scramble random memory
}
```



Dangling pointers

- We **must** eliminate dangling pointers
 - Or type safety is compromised
 - Or memory safety is compromised
 - Or resource safety is compromised
- Eliminated by a combination of rules
 - Distinguish owners from non-owners
 - E.g., `gsl::owner<int*>`
 - Something that holds an owner is an owner
 - Don't forget `malloc()`, etc.
 - Assume raw pointers to be non-owners
 - Catch every attempt for a pointer to “escape” into a scope enclosing its owner's scope
 - `return`, `throw`, out-parameters, long-lived containers, ...



Dangling pointers

- Ensure that no pointer outlives the object it points to

```
void f(X* p)
{
    // ...
    delete p;           // bad: delete non-owner
}

void g()
{
    X* q = new X;       // bad: assign object to non-owner
    f(q);
    // ... do a lot of work here ...
    q->use();           // we never get here
}
```



How to avoid/catch dangling pointers

- Rules (giving pointer safety):
 - Basic rule: no pointer must outlive the object it points to
 - Practical rules
 - Don't transfer to pointer-to-a-local to where it could be accessed by a caller
 - A pointer passed as an argument can be passed back as a result
 - Essential for real-world pointer use
 - A pointer obtained from new can be passed back
 - But we have to remember to eventually delete it

```
int* f(int* p)
{
    int x = 4;
    return &x;           // No! would point to destroyed stack frame
    return new int{7};   // OK: doesn't dangle, but we must "remember" to delete
    return p;           // OK: came from caller
}
```

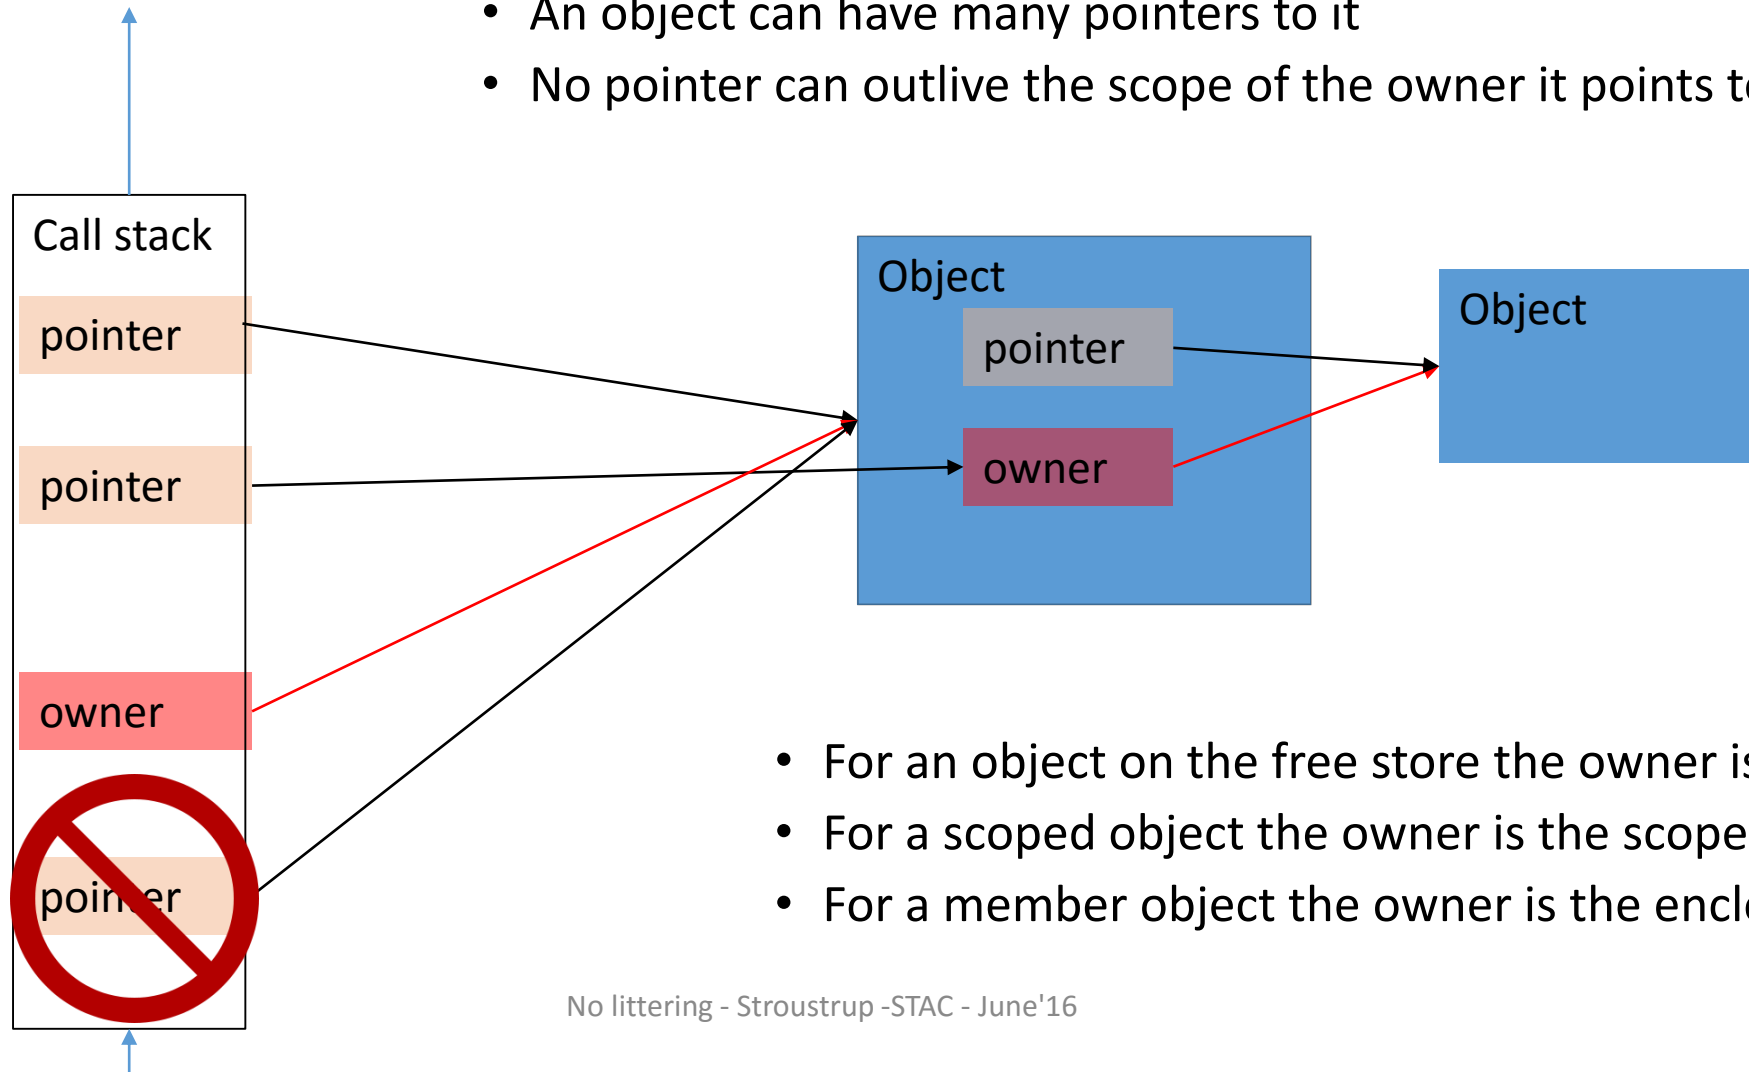

How do we represent ownership?

- Mark an owning **T***: **gsl::owner<T*>**
 - Initial idea (2005 and before)
 - Yet another kind of “smart pointer”
 - **owner<T*>** would hold a **T*** and an “owner bit”
 - Costly: bit manipulation
 - Not ABI compatible
 - Not C compatible
 - Finds errors too late (at run time)
 - So **gsl::owner**
 - Is a handle for static analysis
 - Is documentation
 - Is not a type with it’s own operations
 - Incurs no run-time cost (time or space)
 - Is ABI compatible
 - **template<typename T> using owner = T;**

GSL is our Guidelines Support Library

Owners and pointers

- Every object has one owner
- An object can have many pointers to it
- No pointer can outlive the scope of the owner it points to



- For an object on the free store the owner is a pointer
- For a scoped object the owner is the scope
- For a member object the owner is the enclosing object

How do we manage ownership?

- High-level: Use an ownership abstraction
 - Simple and preferred
 - E.g., `unique_ptr`, `shared_ptr`, `vector`, and `map`
- Low-level: mark owning pointers **owner**
 - An **owner** must be **deleted** or passed to another **owner**
 - A non-**owner** may not be **deleted**
 - This is essential in places but does not scale
 - Applies to both pointers and references



“...And that, in simple terms, is what’s wrong with your software design.”

How do we manage ownership?

- **owner** is intended to simplify static analysis
 - Necessary inside ownership abstractions
 - **owners** in application code is a sign of a problem
 - Usually, C-style interfaces
 - “Lots of annotations” doesn’t scale
 - Becomes a source of errors

GSL: owner<T>

- How do we implement ownership abstractions?

```
template<SemiRegular T>
```

```
class vector {
```

```
public:
```

```
    // ...
```

```
private:
```

```
    owner<T*> elem;
```

```
    // the anchors the allocated memory
```

```
    T* space;
```

```
    // just a position indicator
```


```
    T* end;
```

```
    // just a position indicator
```

```
    // ...
```

```
};
```

- owner<T*>** is just an alias for **T***



GSL: owner<T>

- How about code we cannot change?
 - ABI stability

```
void foo(owner<int*>);      // foo requires an owner
```

```
void f(owner<int*> p, int* q, owner<int*> p2, int* q2)
{
    foo(p);                // OK: transfer ownership
    foo(q);                // bad: q is not an owner
    delete p2;             // necessary
    delete q2;             // bad: not an owner
}
```

- A static analysis tool can tell us where our code mishandles ownership


Our solution: A cocktail of techniques

- Not a single neat miracle cure
 - Rules (from the “Core C++ Guidelines”)
 - Statically enforced
 - Libraries (STL, GSL)
 - So that we don’t have to directly use the messy parts of C++
 - Reliance on the type system
 - The compiler is your friend
 - Static analysis
 - To extend the type system
- None of those techniques is sufficient by itself
- Enforces basic ISO C++ language rules
- Not just for C++
 - But the “cocktail” relies on much of C++



Details (aka engineering)

- Invention is 1% inspiration and 99% perspiration
- The simple lifetime and ownership model needs to be enforced by many dozens of detailed checks
 - Be comprehensive
 - Minimize false positives
 - Scale to industrial programs
 - Fast analysis is essential – local analysis only
 - Allow for gradual adoption
- Provide coherent toolsets for all platforms



“Static” is not quite as flexible as “dynamic”

- Classify pointers according to lifetime

```
int glob = 666;
```

```
vector<int*> f(int* p)
```

```
{
```

```
    int x = 4;
```

```
    int* q = new int{7};
```

```
    vector<int*> res = {p, &x, q, &glob};
```

```
    return res;
```

```
}
```

```
// ignore ownership for now
```

```
// potentially bad: mix lifetimes
```

```
// Bad: return { unknown, &local, free store, &global }
```

- Don't mix different lifetimes in an array (overly conservative?)
 - If you must, encapsulate
- Don't let return statements mix lifetimes

“Static” is not quite as flexible as “dynamic”

- Classify pointers according to ownership

```
int glob = 666;
```

```
vector<int*> f(int* p)
```

```
{
```

```
    int x = 4;
```

```
    owner<int*> q = new int{7};
```

```
    vector<int*> res = {p, &x, q, &glob};    // potentially bad: mix ownership
```

```
    return res;                            // Bad: return {unknown, &local, &owner, &global}
```

```
}
```

- Don't mix different ownerships in an array
 - If you must, encapsulate
- Don't let different return statements mix ownership

Ownership and pointers

- Owners are a tree
 - Except for **shared_ptr**: a DAG
 - Simple
 - efficient
 - Minimal resource retention
 - No ownership cycles
- Owners can be invalidated
 - Catch simple cases at compile time
 - Use **shared_ptr** and/or **nullptr** checks for not-so-simple cases
- Pointers
 - can only refer to live objects
 - To objects with a live owner
 - To objects “back or to the same level” in a stack
 - can have cycles

Concurrency

- Use scopes and **shared_ptr** to keep threads alive as needed
- A thread is a container (of pointers)
 - The usual rules for containers of pointers apply
 - **std::tread**
 - May or may not outlive its scope
 - Bad
 - we must conservatively assume that it lives forever
 - **gsl::raii_thread**
 - Joins
 - so it is a local container
 - **gsl::detached_thread**
 - Detaches
 - so we must treat it as a non-local container

Owner invalidation

- Some cases are simple

```
void f()
{
    auto p = new int{7};
    delete p;          // invalidate p
    *p = 9;             // bad: must be prevented
}
```

- Such examples can be handled by static analysis
 - Avoid “naked new” and “naked delete”

Owner invalidation

- Some cases are less simple

```
void g(int* q) { *q = 9; }      // looks innocent

void f()
{
    vecor<int> v {7};
    auto q = &v[0];
    std::thread t {g,q};
    t.detach();                 // often a bad idea
}

// the thread may outlive v
```

- Such examples can be handled by static analysis
 - Avoid unscoped threads
 - In an emergency, use **shared_ptr** to defeat “false positives”

Owner invalidation

- Some cases are less simple

```
void g(int* q) { *q = 9; }      // looks innocent
```

```
auto f()  
{  
    vecor<int> v {7};  
    auto q = &v[0];  
    return make_shared(thread,g,q); // bad  
}
```

- Such examples can be handled by static analysis

Why not “just use smart pointers”?

- Complexity and (sometimes) cost
 - E.g., different versions of functions for different kinds of pointers
- Use only when you need to manipulate ownership
 - **unique_ptr** for unique ownership
 - guard against exceptions
 - Return pointer-to-base in OOP
 - **shared_ptr** for shared ownership
 - For cases where you can't identify a single owner
 - **Not** for guarding against exceptions
 - **Not** for returning objects from the free store
 - More expensive than raw pointers – use counts
 - Can lead to need for **weak_ptrs**
 - Can lead to “GC delays”
- Remember
 - Local variables (e.g., handles)
 - Move semantics

Static analysis (integrated)

source1-01-05-2016.cpp - Microsoft Visual Studio

FILE EDIT VIEW PROJECT BUILD DEBUG TEAM WINDOW HELP

source1-01-05-2016.cpp

```
79 vector<int*> tst43(int* p)
80 {
81     vector<gsl::owner<int*>> v{ p, new int{ 7 } }; // mixed ownership (bad, IMO)
82     ff42(v); // OK
83     return v; // bad: ownership lost
84 }
85
86 vector<int*> tst44(int* p)
87 {
88     int x = 9;
89     return{ &x, p }; // mixed lifetime, and returning a pointer to local
90 }
91
92 vector<int*> tst45(int* p)
93 {
94     return{ p, new int{ 7 } }; // bad, mixed ownership (n
95 }
96
97 vector<gsl::owner<int*>> tst46(int* p)
98 {
99     return{ p, new int{ 7 } }; // bad, mixed ownership (n
100 }
101
102 struct S42 { int* p; int* q; };
103
104 void ff42(S42&) {}
105
106 S42 tst422(int* p)
107 {
108     S42 v{ p, new int{ 7 } }; // mixed ownership (bad, IMO)
109     ff42(v); // OK, if we allowed the formation of v
110     // but would it be better to catch the formation of the vector (and equivalent array or struct)
111     return v; // bad: ownership lost
112 }
113
114 S42 tst442(int* p)
115 {
116     int x = 9;
```

Code Analysis

Analyze Sort Settings Search

All Projects (44) All Categories (44)

Advanced Filters

C26423

The allocation was not directly assigned to an owner.

c:\repos\source1-01-05-2016.cpp (Line 108)

Miscellaneous

Warning

Actions

C26424

Package Manager Console Output

Ready

No littering - Stroustrup -STAC - June'16

Ln 94 Col 36 Ch 32 INS

Dangling pointer summary

- Simple:
 - Never let a “pointer” escape to where it can refer to its object after that object is destroyed
- It’s not just pointers
 - All ways of “escaping”
 - **return**, **throw**, place in long-lived container, threads, ...
 - Same for containers of pointers
 - E.g. **vector<int*>**, **unique_ptr<int>**, **threads**, iterators, built-in arrays, ...
 - Same for references
- We need a formal paper/proof
- We need to demonstrate scaling
 - 1M line code bases

Other problems

- Other ways of breaking the type system
 - Unions: use variant
 - Casts: don't use them outside abstractions
 - ...
- Other ways of misusing pointers
 - Range errors: use **GSL::span<T>**
 - **nullptr** dereferencing: use **GSL::not_null<T>**
- Wasteful ways of addressing pointer problems
 - Misuse of smart pointers
- ...
- “Just test everywhere at run time” is **not** an acceptable answer
 - We want comprehensive guidelines



GSL::span<T>

- Common interface style

- major source of bugs

```
void f(int* p, int n)           // what is n? (How would a tool know?)
{
    p[7] = 9;                   // OK?
    for (int i=0; i<n; ++i) p[i] = 7; // OK?
}
```

- Better

```
void f(span<int> a)
{
    a[7] = 9;                   // OK? Checkable against a.size()
    for (int& x : a) x = 7;      // OK
}
```

GSL::span<T>

- Common style

```
void f(int* p, int n);  
int a[100];  
// ...  
f(a,100);  
f(a,1000);    // likely disaster
```

- “Make simple things simple”

- Simpler than “old style”
- Shorter
- At least as fast

- Better

```
void f(span<int> a)  
int a[100];  
// ...  
f(span<int>{a});  
f(a);  
f({a,1000}); // easily checkable
```

nullptr problems

- Mixing **nullptr** and pointers to objects
 - Causes confusion
 - Requires (systematic) checking

- Caller

```
void f(char*);
```

```
f(nullptr);           // OK?
```

- Implementer

```
void f(char* p)
```


```
{
```

```
    if (p==nullptr)    // necessary?
```

```
    // ...
```

```
}
```

- Can you trust the documentation?
- Compilers don't read manuals, or comments
- Complexity, errors, and/or run-time cost



GSL::not_null<T>

- Caller

```
void f(not_null<char*>);
```

```
f(nullptr);    // Obvious error: caught by static analysis
```

```
char* p = nullptr;
```

```
f(p);          // Constructor for not_null can catch the error
```

- Implementer

```
void f(not_null<char*> p)
```

```
{
```

```
    // if (p==nullptr) // not necessary
```

```
    // ...
```

```
}
```

GSL::not_null<T>

- **not_null<T>**
 - A simple, small class
 - Should it be an alias like **owner**?
 - **not_null<T*>** is **T*** except that it cannot hold **nullptr**
 - Can be used as input to analyzers
 - Minimize run-time checking
 - Checking can be “debug only”
 - For any **T** that can be compared to **nullptr**

To summarize

- Type and resource safety:
 - RAI (scoped objects with constructors and destructors)
 - No dangling pointers
 - No leaks (track ownership pointers)
 - Eliminate range errors
 - Eliminate nullptr dereference
- That done, we attack other sources of problems
 - Logic errors
 - Performance bugs
 - Maintenance hazards
 - Verbosity
 - ...



Current state: the game is changing dramatically

- Papers


- B. Stroustrup, H. Sutter, G. Dos Reis: A brief introduction to C++'s model for type- and resource-safety.
- H. Sutter and N. MacIntosh: Preventing Leaks and Dangling
- T. Ramananandro, G. Dos Reis, X Leroy: A Mechanized Semantics for C++ Object Construction and Destruction, with Applications to Resource Management

- Code (MIT license)

- <https://github.com/isocpp/CppCoreGuidelines>
- <https://github.com/microsoft/gsl>
- Static analysis: experimental versions available (Microsoft)

- Videos

- B. Stroustrup: : Writing Good C++ 14
- H. Sutter: Writing good C++ 14 By Default
- G. Dos Reis: Contracts for Dependable C++
- N. MacIntosh: Static analysis and C++: more than lint
- N. MacIntosh: A few good types: Evolving array_view and string_view for safe C++ code





We are not unambitious (rough seas ahead)

- Type and resource safety
 - No leaks
 - No dangling pointers
 - No bad accesses
 - No range errors
 - No use of uninitialized objects
 - No misuse of
 - Casts
 - Unions
- We think we can do it
 - At scale
 - 4+ million C++ Programmers, N billion lines of code
 - Zero-overhead principle



Questions?





- Type- and Resource-safe C++
 - No garbage collector (because there is no garbage to collect)
 - No runtime overheads (Except necessary range checks)
 - No new limits on expressibility
 - ISO C++
 - Simpler code



C++ Information

- The C++ Foundation: www.isocpp.org
 - Standards information, articles, user-group information
- Bjarne Stroustrup: www.stroustrup.com
 - Publication list, C++ libraries, FAQs, etc.
 - *A Tour of C++*: All of C++ in 180 pages
 - *The C++ Programming Language* (4th edition): All of C++ in 1,300 pages
 - *Programming: Principles and Practice using C++* (2nd edition): A textbook
- The ISO C++ Standards Committee: www.open-std.org/jtc1/sc22/wg21/
 - All committee documents (incl. proposals)



Videos

- Cppcon: <https://www.youtube.com/user/CppCon> 2014, 2015
- Going Native: <http://channel9.msdn.com/Events/GoingNative/> 2012, 2013

Guidelines: <https://github.com/isocpp/CppCoreGuidelines>

