

Best Practices

Rainer Grimm
Training, Coaching, and
Technology Consulting
www.ModernesCpp.de

Best Practices

General

Multithreading

Parallel

Memory Model

Best Practices

General

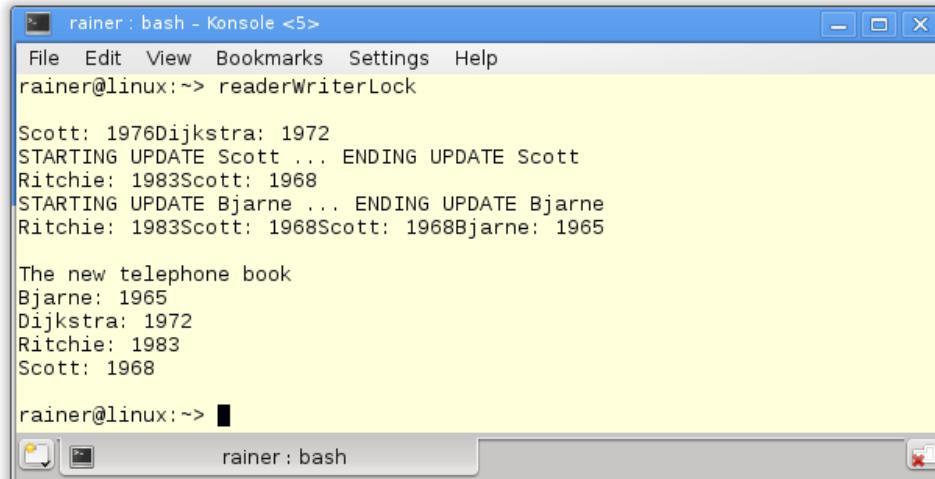
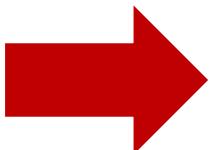
Multithreading

Memory Model

Code Reviews

```
map<string,int> teleBook{{"Dijkstra", 1972},  
                         {"Scott", 1976}, {"Ritchie", 1983}};  
  
shared_timed_mutex teleBookMutex;  
  
void addToTeleBook(const string& na, int tele){  
    lock_guard<shared_timed_mutex> writerLock(teleBookMutex);  
    cout << "\nSTARTING UPDATE " << na;  
    this_thread::sleep_for(chrono::milliseconds(500));  
    teleBook[na]= tele;  
    cout << " ... ENDING UPDATE " << na << endl;  
}  
  
void printNumber(const string& na){  
    shared_lock<shared_timed_mutex> readerLock(teleBookMutex);  
    cout << na << ":" << teleBook[na] << endl;  
}
```

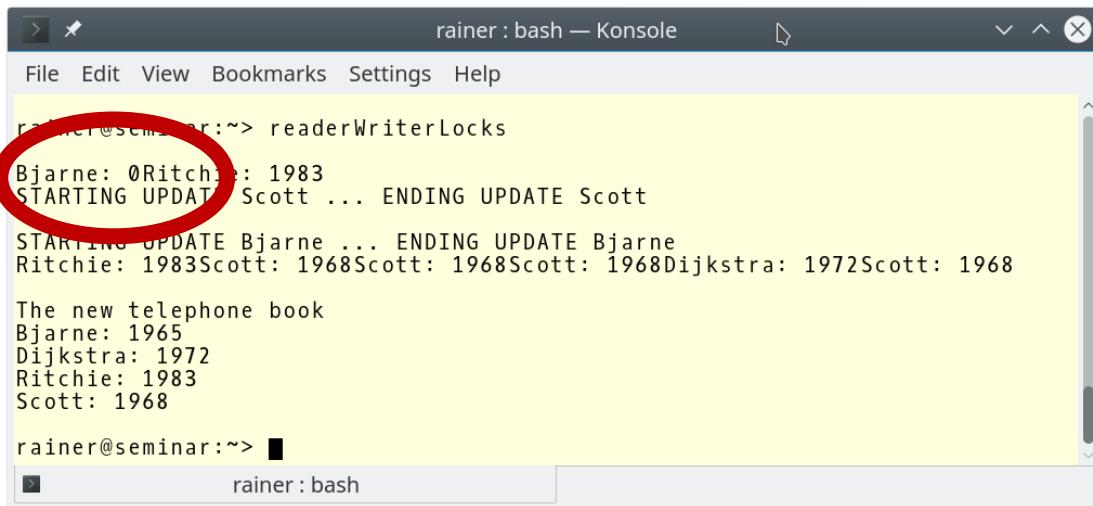
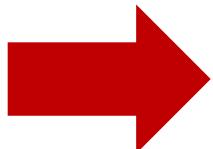
```
thread reader1([]{ printNumber("Scott"); });  
thread reader2([]{ printNumber("Ritchie"); });  
thread w1([]{ addToTeleBook("Scott",1968); });  
thread reader3([]{ printNumber("Dijkstra"); });  
thread reader4([]{ printNumber("Scott"); });  
thread w2([]{ addToTeleBook("Bjarne",1965); });  
thread reader5([]{ printNumber("Scott"); });  
thread reader6([]{ printNumber("Ritchie"); });  
thread reader7([]{ printNumber("Scott"); });  
thread reader8([]{ printNumber("Bjarne"); });  
  
reader1.join(), reader2.join();  
reader3.join(), reader4.join();  
reader5.join(), reader6.join();  
reader7.join(), reader8.join();  
w1.join(), w2.join();  
  
cout << "\nThe new telephone book" << endl;  
for (auto teleIt: teleBook){  
    cout << teleIt.first << ":" << teleIt.second << endl;  
}
```



Code Reviews

```
map<string,int> teleBook{{"Dijkstra", 1972},  
                         {"Scott", 1976}, {"Ritchie", 1983}};  
  
shared_timed_mutex teleBookMutex;  
  
void addToTeleBook(const string& na, int tele){  
    lock_guard<shared_timed_mutex> writerLock(teleBookMutex);  
    cout << "\nSTARTING UPDATE " << na;  
    this_thread::sleep_for(chrono::milliseconds(500));  
    teleBook[na]= tele;  
    cout << " ... ENDING UPDATE " << na << endl;  
}  
  
void printNumber(const string& na){  
    shared_lock<shared_timed_mutex> readerLock(teleBookMutex);  
    cout << na << ":" << teleBook[na] << endl;  
}
```

```
thread reader1([]{ printNumber("Scott"); });  
thread reader2([]{ printNumber("Ritchie"); });  
thread w1([]{ addToTeleBook("Scott",1968); });  
thread reader3([]{ printNumber("Dijkstra"); });  
thread reader4([]{ printNumber("Scott"); });  
thread w2([]{ addToTeleBook("Bjarne",1965); });  
thread reader5([]{ printNumber("Scott"); });  
thread reader6([]{ printNumber("Ritchie"); });  
thread reader7([]{ printNumber("Scott"); });  
thread reader8([]{ printNumber("Bjarne"); });  
  
reader1.join(), reader2.join();  
reader3.join(), reader4.join();  
reader5.join(), reader6.join();  
reader7.join(), reader8.join();  
w1.join(), w2.join();  
  
cout << "\nThe new telephone book" << endl;  
for (auto teleIt: teleBook){  
    cout << teleIt.first << ":" << teleIt.second << endl;  
}
```



Minimize Sharing

- Summation of a vector with 100 000 000 elements

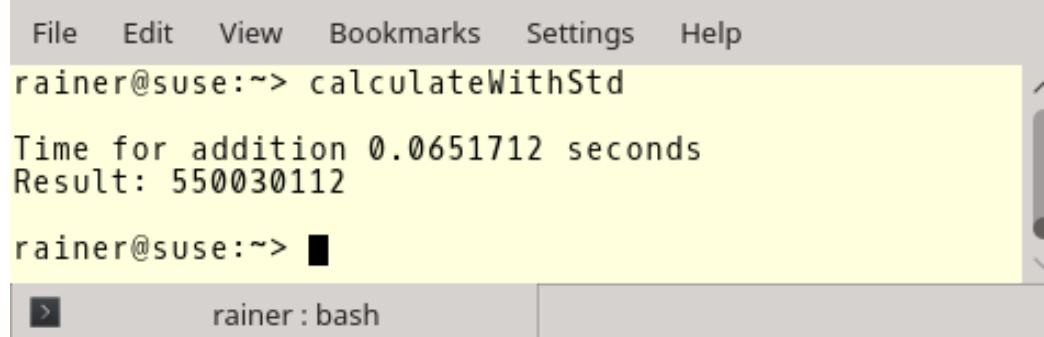
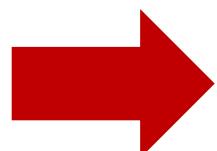
```
constexpr long long size = 100000000;  
...  
// random values  
std::vector<int> randValues;  
randValues.reserve(size);  
  
std::random_device seed;  
std::mt19937 engine(seed());  
std::uniform_int_distribution<> uniformDist(1, 10);  
for (long long i = 0 ; i < size ; ++i)  
    randValues.push_back(uniformDist(engine));  
  
...  
// calculate sum  
...
```

Minimize Sharing

- Single-threaded in two variations

```
unsigned long long sum {};
for (auto n: randValues) sum += n;

const unsigned long long sum = accumulate(randValues.begin(),
                                         randValues.end(), 0ll);
```



A screenshot of a terminal window titled "rainer : bash". The window shows the following text:

```
File Edit View Bookmarks Settings Help
rainer@suse:~> calculateWithStd
Time for addition 0.0651712 seconds
Result: 550030112
rainer@suse:~> █
▶ rainer : bash
```

The terminal window has a standard Linux-style interface with a menu bar at the top. The command `calculateWithStd` was run, followed by its output: the time taken for the addition (0.0651712 seconds) and the resulting sum (550030112). A red arrow on the left points to the terminal window, highlighting the output.

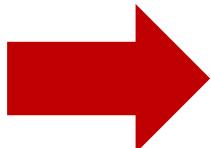
Minimize Sharing

- Four threads with a shared summation variable

```
void sumUp(unsigned long long& sum, const vector<int>& val,
           unsigned long long beg, unsigned long long end){
    for (auto it = beg; it < end; ++it){
        lock_guard<mutex> myLock(myMutex);
        sum += val[it];
    }
}

...
unsigned long long sum{};

thread t1(sumUp, ref(sum), ref(randValues), 0, fir);
thread t2(sumUp, ref(sum), ref(randValues), fir, sec);
thread t3(sumUp, ref(sum), ref(randValues), sec, thi);
thread t4(sumUp, ref(sum), ref(randValues), thi, fou);
```



The terminal window shows the execution of a C++ program named `calculateWithLock`. The program calculates the sum of elements in a vector using four threads. The output indicates that the addition took 3.3389 seconds and the result is 549961505. The terminal prompt is `rainer@suse:~>`.

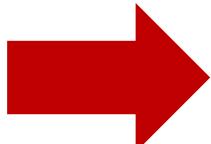
```
File Edit View Bookmarks Settings Help
rainer@suse:~> calculateWithLock
Time for addition 3.3389 seconds
Result: 549961505
rainer@suse:~> □
> rainer : bash
```

Minimize Sharing

- Four threads with a shared, atomic summation variable

```
void sumUp(atomic<unsigned long long>& sum, const vector<int>& val,
            unsigned long long beg, unsigned long long end){
    for (auto it = beg; it < end; ++it){
        sum += val[it];
    }
}

void sumUp(atomic<unsigned long long>& sum, const vector<int>& val,
            unsigned long long beg, unsigned long long end){
    for (auto it = beg; it < end; ++it){
        sum.fetch_add(val[it], memory_order_relaxed);
    }
}
```



```
File Edit View Bookmarks Settings Help
rainer@suse:~> calculateWithAtomic
sum.is_lock_free(): true
Time for addition 1.33837 seconds
Result: 549992025
Time for addition 1.34625 seconds
Result: 549992025
rainer@suse:~> █
> rainer : bash
```

Minimize Sharing

- Four threads with a local summation variable

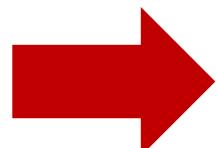
```
void sumUp(unsigned long long& sum, const vector<int>& val,
           unsigned long long beg, unsigned long long end){
    unsigned long long tmpSum{};
    for (auto i = beg; i < end; ++i){
        tmpSum += val[i];
    }
    lock_guard<mutex> lockGuard(myMutex);
    sum += tmpSum;
}
```

File Edit View Bookmarks Settings Help

rainer@suse:~> localVariable

Time for addition 0.0284271 seconds
Result: 549996948

rainer@suse:~> █



> rainer:bash

Minimize Sharing

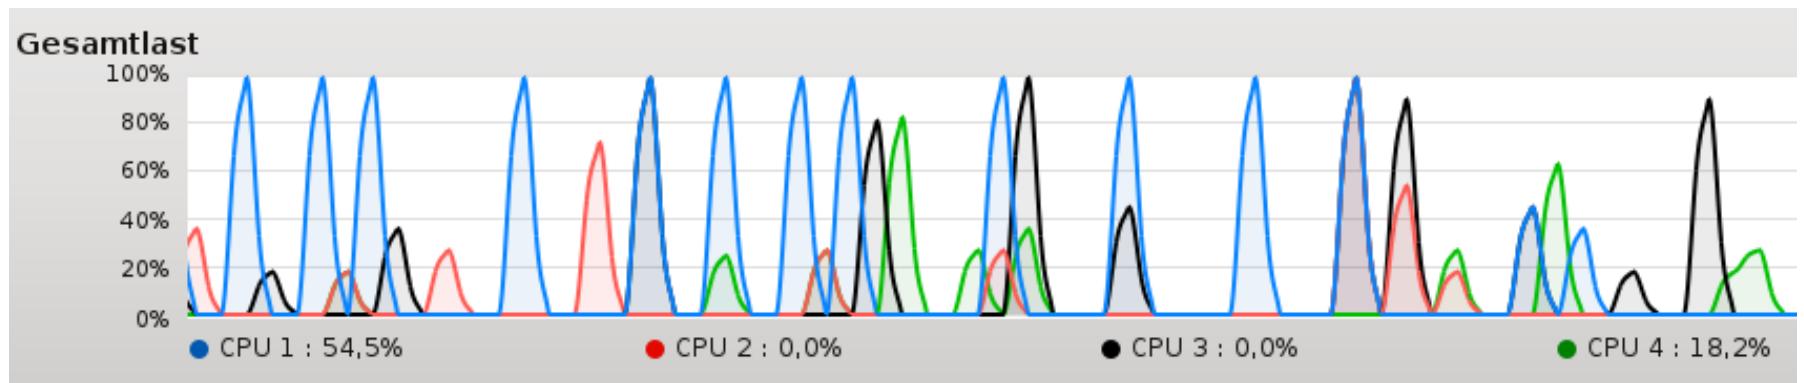
- The results

Single-threaded	4 threads with a lock	4 threads with an atomic	4 threads with an local variable
0.07 sec	3.34 sec	1.34 sec	0.03 sec

All Fine?

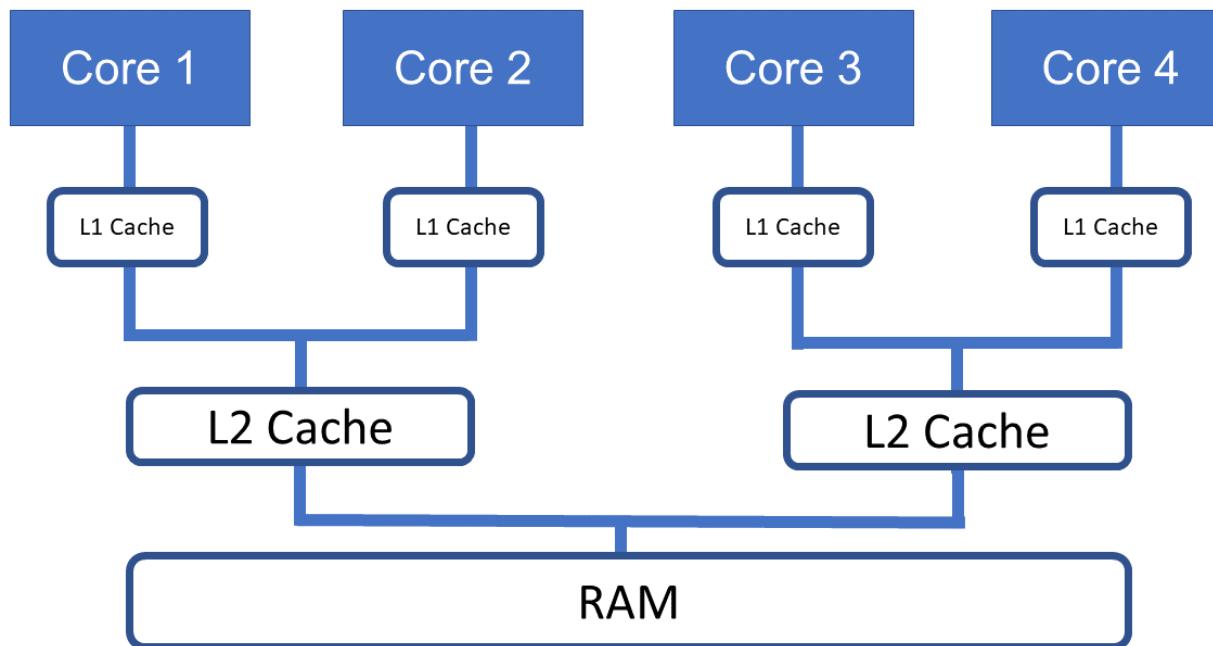
Minimize Sharing

- The CPU utilisation



Minimize Sharing

- The memory wall



The RAM is the bottleneck.

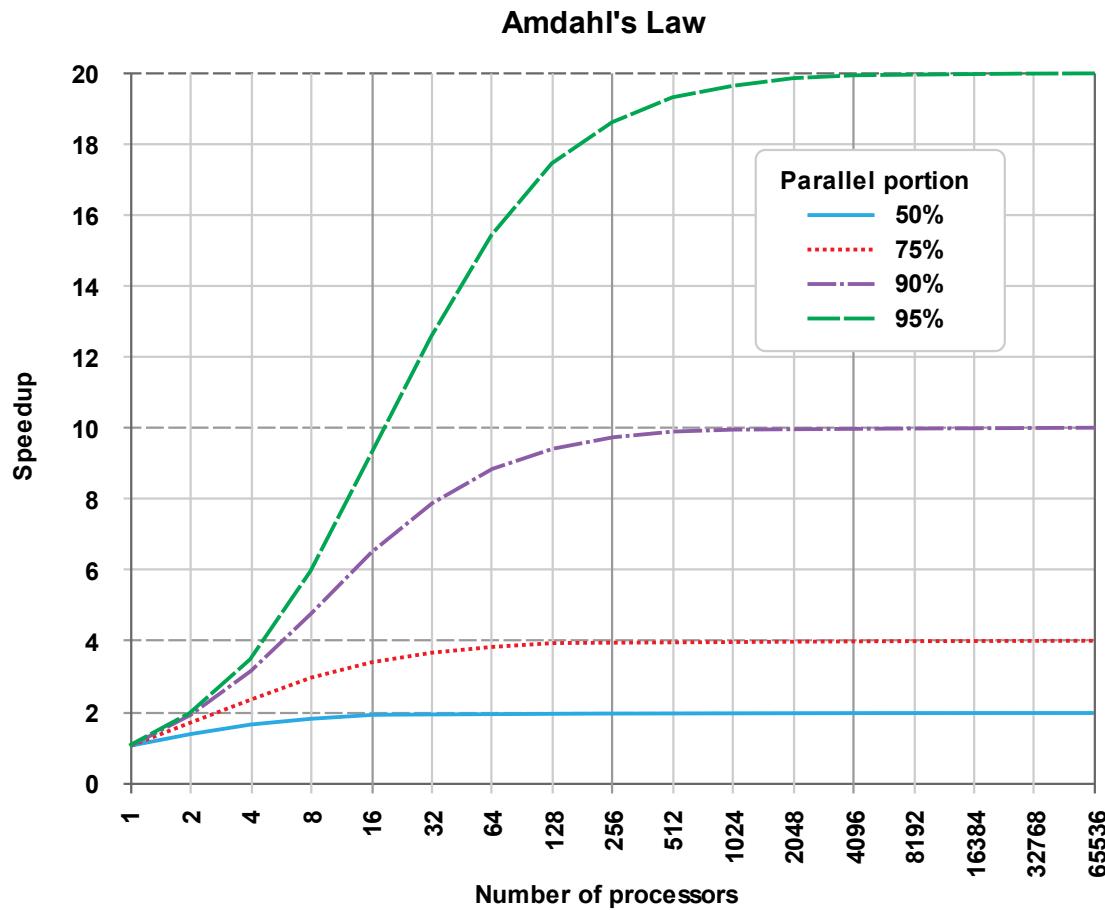
Minimize Waiting (Amdahl's Law)

$$\frac{1}{1 - p}$$

p: Parallel Code

$$p = 0.5 \rightarrow 2$$

Minimize Waiting (Amdahl's Law)



By Daniels220 at English Wikipedia, CC BY-SA 3.0, <https://commons.wikimedia.org/w/index.php?curid=6678551>

Use Code Analysis Tools (ThreadSanitizer)

- Detects Data Races
- Memory- and Performance overhead: 10x

```
#include <thread>

int main() {

    int globalVar{};

    std::thread t1([&globalVar]{ ++globalVar; });
    std::thread t2([&globalVar]{ ++globalVar; });
    t1.join(), t2.join();

}
```

Use Code Analysis Tools (ThreadSanitizer)

```
g++ dataRace.cpp -fsanitize=thread -pthread -g -o dataRace
```

```
File Edit View Bookmarks Settings Help
rainer@suse:~> dataRace
=====
WARNING: ThreadSanitizer: data race (pid=6764)
  Read of size 4 at 0x7fff031ca3bc by thread T2:
#0 operator() /home/rainer/dataRace.cpp:10 (dataRace+0x000000400f01)
#1 __invoke<> /usr/local/include/c++/6.3.0/functional:1391 (dataRace+0x000000401b3d)
#2 operator() /usr/local/include/c++/6.3.0/functional:1380 (dataRace+0x000000401a51)
#3 __run /usr/local/include/c++/6.3.0/thread:196 (dataRace+0x0000004019bc)
#4 execute_native_thread_routine ../../../../../libstdc++-v3/src/c++11/thread.cc:83 (libstdc++.so.6+0x00000000c11be)

  Previous write of size 4 at 0x7fff031ca3bc by thread T1:
#0 operator() /home/rainer/dataRace.cpp:9 (dataRace+0x000000400eb9)
#1 __invoke<> /usr/local/include/c++/6.3.0/functional:1391 (dataRace+0x000000401be7)
#2 operator() /usr/local/include/c++/6.3.0/functional:1380 (dataRace+0x000000401a8b)
#3 __run /usr/local/include/c++/6.3.0/thread:196 (dataRace+0x000000401a06)
#4 execute_native_thread_routine ../../../../../libstdc++-v3/src/c++11/thread.cc:83 (libstdc++.so.6+0x00000000c11be)

  Location is stack of main thread.

  Thread T2 (tid=6767, running) created by main thread at:
#0 pthread_create ../../../../../libasanizer/tsan/tsan_interceptors.cc:876 (libtsan.so.0+0x00000002aaed)
#1 __gthread_create /home/rainer/languages/C++/gcc-6.3.0/x86_64-pc-linux-gnu/libstdc++-v3/include/x86_64-pc-linux-gnu/bits/gthr-default.h:662 (libstdc++.so.6+0x0000000c14b4)
#2 std::thread::_M_start_thread(std::unique_ptr<std::thread::State, std::default_delete<std::thread::Stat e> >, void (*)()) ../../../../../libstdc++-v3/src/c++11/thread.cc:163 (libstdc++.so.6+0x0000000c14b4)
#3 main /home/rainer/dataRace.cpp:10 (dataRace+0x000000400f97)

  Thread T1 (tid=6766, finished) created by main thread at:
#0 pthread_create ../../../../../libasanizer/tsan/tsan_interceptors.cc:876 (libtsan.so.0+0x00000002aaed)
#1 __gthread_create /home/rainer/languages/C++/gcc-6.3.0/x86_64-pc-linux-gnu/libstdc++-v3/include/x86_64-pc-linux-gnu/bits/gthr-default.h:662 (libstdc++.so.6+0x0000000c14b4)
#2 std::thread::_M_start_thread(std::unique_ptr<std::thread::State, std::default_delete<std::thread::Stat e> >, void (*)()) ../../../../../libstdc++-v3/src/c++11/thread.cc:163 (libstdc++.so.6+0x0000000c14b4)
#3 main /home/rainer/dataRace.cpp:9 (dataRace+0x000000400f70)

SUMMARY: ThreadSanitizer: data race /home/rainer/dataRace.cpp:10 in operator()

ThreadSanitizer: reported 1 warnings
rainer@suse:~>
```

Use Code Analysis Tools (ThreadSanitizer)

```
bool dataReady= false;  
  
std::mutex mutex_;  
std::condition_variable condVar1;  
std::condition_variable condVar2;  
  
int counter=0;  
int COUNTLIMIT=50;  
  
void setTrue(){  
  
    while(counter <= COUNTLIMIT){  
  
        std::unique_lock<std::mutex> lck(mutex_);  
        condVar1.wait(lck=[]{return dataReady == false;});  
        dataReady= true;  
        ++counter;  
        std::cout << dataReady << std::endl;  
        condVar2.notify_one();  
    }  
}
```

```
void setFalse(){  
  
    while(counter < COUNTLIMIT){  
  
        std::unique_lock<std::mutex> lck(mutex_);  
        condVar2.wait(lck=[]{return dataReady == true;});  
        dataReady= false;  
        std::cout << dataReady << std::endl;  
        condVar1.notify_one();  
    }  
}  
  
int main(){  
  
    std::cout << std::boolalpha << std::endl;  
  
    std::cout << "Begin: " << dataReady << std::endl;  
  
    std::thread t1(setTrue);  
    std::thread t2(setFalse);  
  
    t1.join();  
    t2.join();  
  
    dataReady= false;  
    std::cout << "End: " << dataReady << std::endl;  
    std::cout << std::endl;  
}
```

Use Code Analysis Tools (ThreadSanitizer)

```
File Edit View Bookmarks Settings Help
rainer@linux: ~ conditionvariablepingpong
.
Begin: false
true
=====
WARNING: ThreadSanitizer: data race (pid=18139)
  Read of size 4 at 0x000000004350 by thread T2:
#0 setFalse() /home/rainer/conditionVariablePingPong.cpp:30 (conditionVariablePingPong+0x000000401818)
#1 void std::__Bind_simple<void (*())()>::__M_invoke<(std::__Index_tuple<>) /usr/include/c++/6/functional:1400
#2 std::__Bind_simple<void (*())()>::operator()() /usr/include/c++/6/functional:1389 (conditionVariablePingPong+0x000000401818)
#3 std::thread::__State_impl<std::__Bind_simple<void (*())()>::__M_run() /usr/include/c++/6/thread:196 (conditionVariablePingPong+0x000000401818)
#4 <null> <null> (libstdc++.so.6+0x0000000c22de)

  Previous write of size 4 at 0x000000604350 by thread T1 (mutexes: write M11):
#0 setTrue() /home/rainer/conditionVariablePingPong.cpp:21 (conditionVariablePingPong+0x00000040173d)
#1 void std::__Bind_simple<void (*())()>::__M_invoke<(std::__Index_tuple<>) /usr/include/c++/6/functional:1400
#2 std::__Bind_simple<void (*())()>::operator()() /usr/include/c++/6/functional:1389 (conditionVariablePingPong+0x00000040173d)
#3 std::thread::__State_impl<std::__Bind_simple<void (*())()>::__M_run() /usr/include/c++/6/thread:196 (conditionVariablePingPong+0x00000040173d)
#4 <null> <null> (libstdc++.so.6+0x0000000c22de)

Location is global 'counter' of size 4 at 0x000000604350 (conditionVariablePingPong+0x000000604350)

Mutex M11 (0x0000006042a0) created at:
#0 pthread_mutex_lock <null> (libtsan.so.0+0x00000003bc0f)
#1 __gthread_mutex_lock /usr/include/c++/6/x86_64-suse-linux/bits/gthr-default.h:748 (conditionVariablePingPong+0x000000401be0)
#2 std::mutex::lock() /usr/include/c++/6/bits/std_mutex.h:103 (conditionVariablePingPong+0x000000401be0)
#3 std::unique_lock<std::mutex>::lock() /usr/include/c++/6/bits/std_mutex.h:267 (conditionVariablePingPong+0x000000401be0)
#4 std::unique_lock<std::mutex>::unique_lock(std::mutex&) /usr/include/c++/6/bits/std_mutex.h:197 (conditionVariablePingPong+0x000000401be0)
#5 setTrue() /home/rainer/conditionVariablePingPong.cpp:18 (conditionVariablePingPong+0x0000004016f4)
#6 void std::__Bind_simple<void (*())()>::__M_invoke<(std::__Index_tuple<>) /usr/include/c++/6/functional:1400
#7 std::__Bind_simple<void (*())()>::operator()() /usr/include/c++/6/functional:1389 (conditionVariablePingPong+0x0000004016f4)
#8 std::thread::__State_impl<std::__Bind_simple<void (*())()>::__M_run() /usr/include/c++/6/thread:196 (conditionVariablePingPong+0x0000004016f4)
#9 <null> <null> (libstdc++.so.6+0x0000000c22de)

Thread T2 (tid=18140, running) created by main thread at:
#0 pthread_create <null> (libtsan.so.0+0x00000002b740)
#1 std::thread::__M_start_thread(std::unique_ptr<std::thread::__State, std::default_delete<std::thread::__State> 5d4)
#2 main /home/rainer/conditionVariablePingPong.cpp:49 (conditionVariablePingPong+0x00000040197c)

Thread T1 (tid=18139, running) created by main thread at:
#0 pthread_create <null> (libtsan.so.0+0x00000002b740)
#1 std::thread::__M_start_thread(std::unique_ptr<std::thread::__State, std::default_delete<std::thread::__State> 5d4)
#2 main /home/rainer/conditionVariablePingPong.cpp:48 (conditionVariablePingPong+0x00000040196b)

SUMMARY: ThreadSanitizer: data race /home/rainer/conditionVariablePingPong.cpp:30 in setFalse()
=====
false
true
false
true
false
```

rainer : bash

Use Code Analysis Tools (CppMem)

CppMem: Interactive C/C++ memory model

Model
standard preferred release_acquire tot relaxed_only

Program
examples/Paper | data_race.c

C Execution

```
// a data race (dr)
int main() {
    int x = 2;
    int y;
    {{ x = 3;
    ||| y = (x==3);
    }};
    return 0;
}
```

2 reset help **2 executions; 1 consistent, not race free**

Computed executions

Display Relations

sb asw dd cd
 rf mo sc lo
 hb vse ithb sw rs hrs dob cad
 unsequenced_races data_races

Display Layout

dot neato_par neato_par_init neato_downwards
 tex

[edit display options](#)

3

Execution candidate no. 2 of 2

previous consistent previous candidate next candidate next consistent 2 goto

Model Predicates

consistent_race_free_execution = **false**
 consistent_execution = **true**
 assumptions = **true**
 well_formed_threads = **true**
 well_formed_rf = **true**
 locks_only_consistent_locks = **true**
 locks_only_consistent_lo = **true**
 consistent_mo = **true**
 sc_accesses_consistent_sc = **true**
 sc_fenced_sc_fences_heeded = **true**
 consistent_hb = **true**
 consistent_rf = **true**
 det_read = **true**
 consistent_non_atomic_rf = **true**
 consistent_atomic_rf = **true**
 coherent_memory_use = **true**
 rmw_atomicity = **true**
 sc_accesses_sc_reads_restricted = **true**
unsequenced_races are **absent**
data_races are **present**
ineterminate_reads are **absent**
locks_only_bad_mutexes are **absent**

a:Wna x=2

sw

rf,sw

dr

sb

sb

sb

sb

sb

4

Files: out.exc, out.dot, out.dsp, out.tex

Use Code Analysis Tools (CppMem)

```
int x = 0, std::atomic<int> y{0};

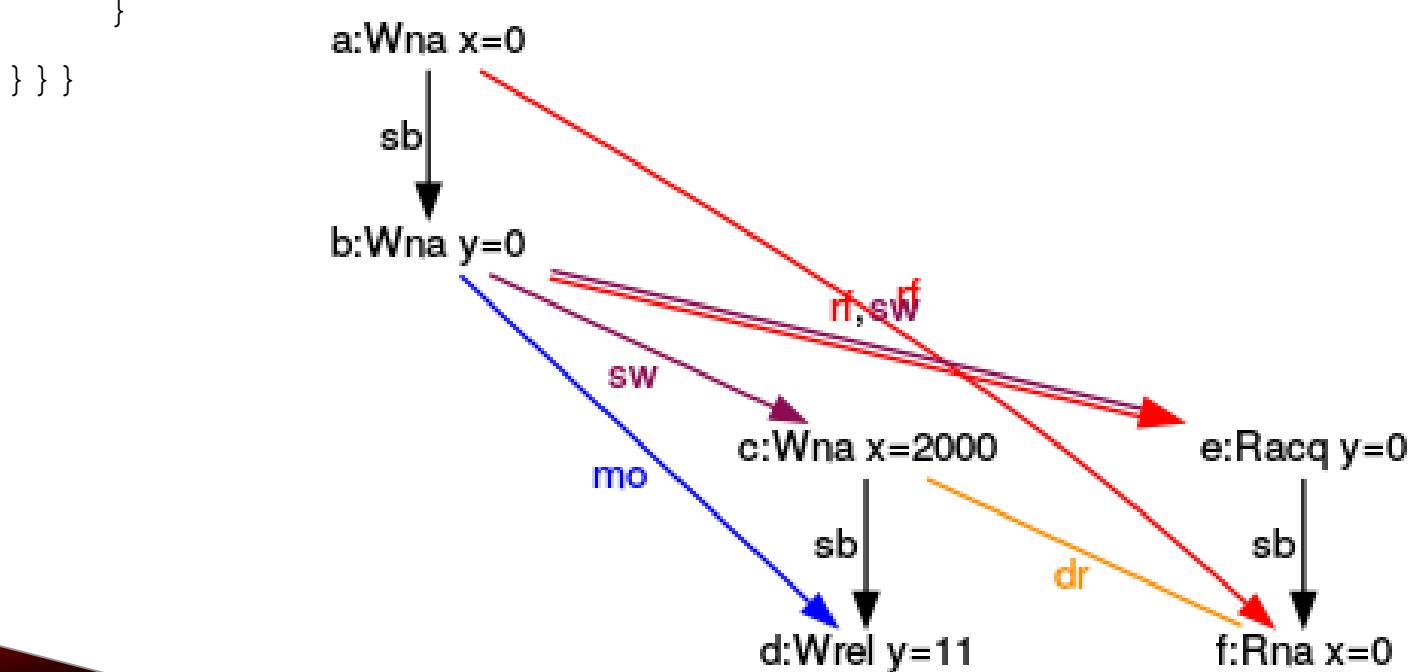
void writing() {
    x = 2000;
    y.store(11, std::memory_order_release);
}

void reading() {
    std::cout << y.load(std::memory_order_acquire) << " ";
    std::cout << x << std::endl;
}

std::thread thread1(writing);
std::thread thread2(reading);
thread1.join(), thread2.join();
```

Use Code Analysis Tools (CppMem)

```
int x = 0, atomic_int y = 0;  
{ {{ {  
    x = 2000;  
    y.store(11, memory_order_release);  
}  
| || | {  
    y.load(memory_order_acquire);  
    x;  
}  
}}}
```



Use Immutable Data

Data Race: At least two threads access a shared variable at the same time. At least one thread tries to modify the it.

Mutable?

Shared?

	no	yes
no	OK	Ok
yes	OK	Data Race

Use Immutable Data

The remaining Problem: initialise the data thread-safe

1. Early Initialisation

```
const int val = 2011;  
thread t1([&val]{ .... });  
thread t2([&val]{ .... });
```

3. `call_once` and `once_flag`

```
void onlyOnceFunc{ .... };  
call_once(onceFlag, onlyOnceFunc);  
thread t3{ onlyOnceFunc() };  
thread t4{ onlyOnceFunc() };
```

2. Constant Expressions

```
constexpr auto doub = 5.1;
```

4. Static variables in a scope

```
void func() {  
    .... static int val 2011; ....  
}  
thread t5{ func() };  
thread t6{ func() };
```

Use Pure Functions

Pure Function

Produce always the same result when given the same arguments

Has no side effect

Don't change the global state of the program

Added Value

- Easier to make correctness proofs
- Refactoring and testing is easier
- Results from previous function calls can be memorised
- **The sequence of function invocations can automatically be changed or the function can automatically be parallelised**

Use Pure Functions

- Function

```
int powFunc(int m, int n){  
    if (n == 0) return 1;  
    return m * powFunc(m, n-1);  
}
```

- Metaprogramming

```
template<int m, int n>  
struct PowMeta{  
    static int const value = m * PowMeta<m, n-1>::value;  
};  
  
template<int m>  
struct PowMeta<m, 0>{  
    static int const value = 1;  
};
```

Use Pure Functions

- `constexpr` Function
 - almost pure functions

```
constexpr int powConst(int m, int n) {  
    int r = 1;  
    for(int k = 1; k <= n; ++k) r *= m;  
    return r;  
}
```

```
auto res = powConst(2, 10);
```

```
auto constexpr res2 = powConst(2, 10);
```

Best Practices

General

Multithreading

Memory Model

Use Tasks instead of Threads

Thread

```
int res;  
thread t([&]{ res = 3 + 4; });  
t.join();  
cout << res << endl;
```

Task

```
auto fut = async([]{ return 3 + 4; });  
cout << fut.get() << endl;
```

Criteria	Thread	Task
Parties Involved	creator thread and child thread	promise and future
Communication	shared variable	communication channel
Thread Creation	obligatory	optional
Synchronisation	join call blocks	get call blocks
Exception in Child	creator thread and child thread die	return value of the promise
Forms of Communication	values	values, notifications, and threads

Use Tasks instead of Threads

C++20: Extend futures will support composition

- **then**: Execute the future if the previous future is done
- **when_any**: Execute the future if any of the previous future is done
- **when_all**: Execute the future if all of the previous futures are done

Use Tasks instead of Condition Variables

Thread 1

```
{  
    lock_guard<mutex> lck(mut);  
    ready = true;  
}  
condVar.notify_one();
```

Thread 2

```
{  
    unique_lock<mutex>lck(mut);  
    condVar.wait(lck, []{ return ready; });  
}
```

prom.set_value();  fut.wait();

Criteria	Condition Variable	Tasks
Critical Region	yes	no
Spurious Wakeup	yes	no
Lost Wakeup	yes	no
Repeated Synchronisation	yes	no

Pack Mutexes in Locks

No Release of the lock

- Mutex

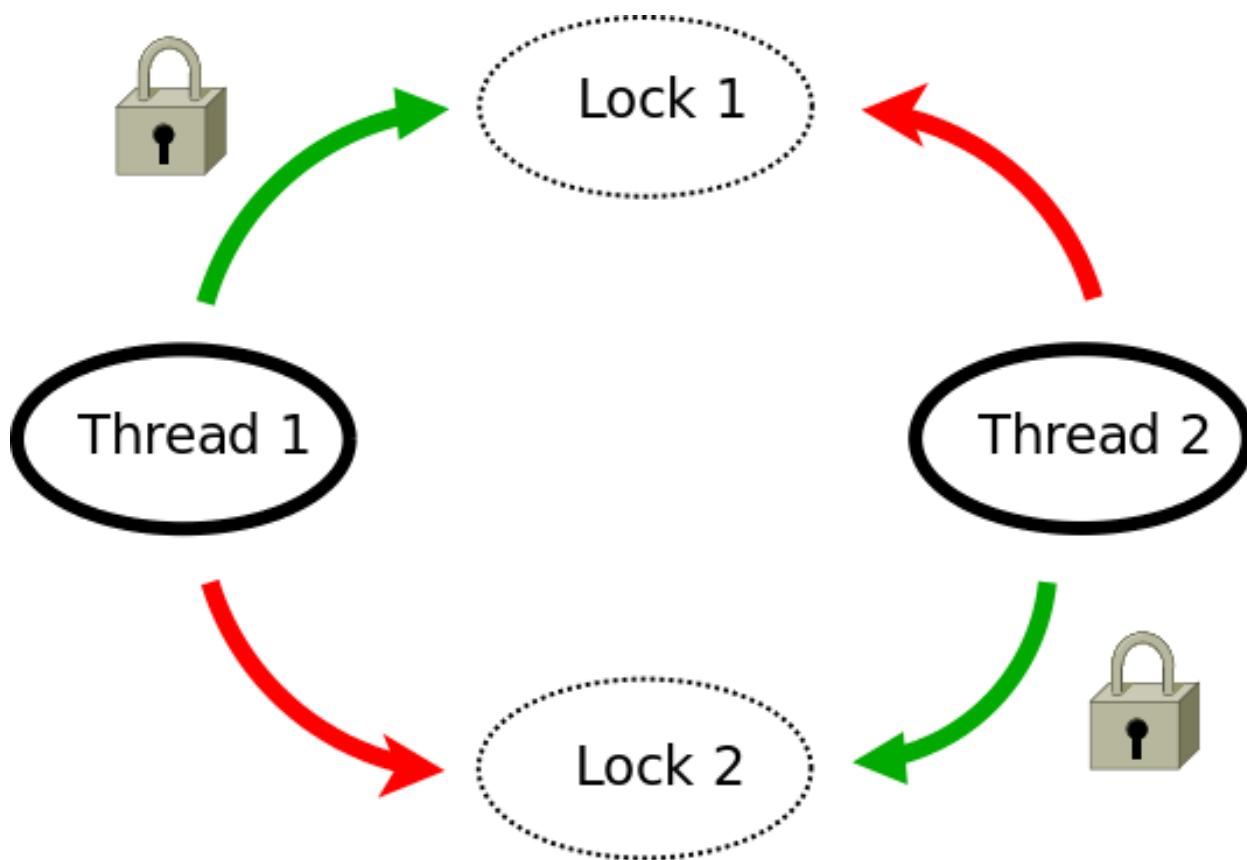
```
mutex m;  
  
{  
    m.lock();  
    shrVar = getVar();  
    m.unlock();  
}
```

- lock_guard

```
mutex m;  
  
{  
    lock_guard<mutex> myLock(m);  
    shaVar = getVar();  
}
```

Pack Mutexes in Locks

Locking of the mutexes is different order



Pack Mutexes in Locks

Atomic lock of the mutex

- unique_lock

```
{  
    unique_lock<mutex> guard1(mut1, defer_lock);  
    unique_lock<mutex> guard2(mut2, defer_lock);  
    lock(guard1, guard2);  
}
```

- scoped_lock (C++17)

```
{  
    std::scoped_lock mut1, mut2;  
}
```

Best Practices

General

Multithreading

Memory Model

Don't Program lock-free



Guide to Threaded Coding

1. Forget what you learned in Kindergarten
(ie stop Sharing)
2. Use Locks
3. Measure
4. Measure
5. Change your Algorithm
6. GOTO 1

Tony Van Eerd

∞. Lock-free

Lock-free coding is the last thing you want to do.

Fedor Pikus

- Writing lock-free programs is hard
- Writing correct lock-free programs is even harder

Safety: off
How not to shoot yourself in the foot with C++ atomics



The ugly side of weakly ordered atomics

Extreme complexity.

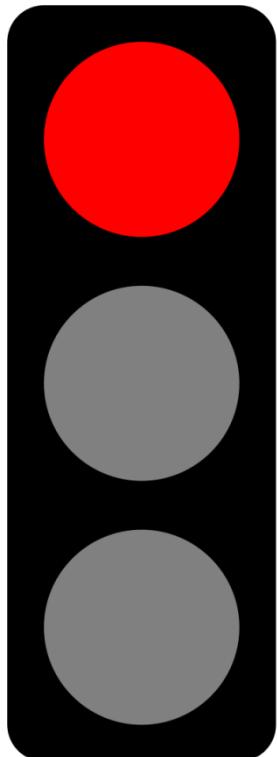
- The rules are not obvious.
- They're often downright surprising.
- And not even well understood.
- The committee still hasn't figured out how to define `memory_order_relaxed`.
... and I'm not even going to talk about `memory_order_consume`.

Hans Böhm

The specification of release-consume ordering is being revised, and the use of `memory_order_consume` is (since C++17) temporarily discouraged.

Don't Program lock-free: ABA

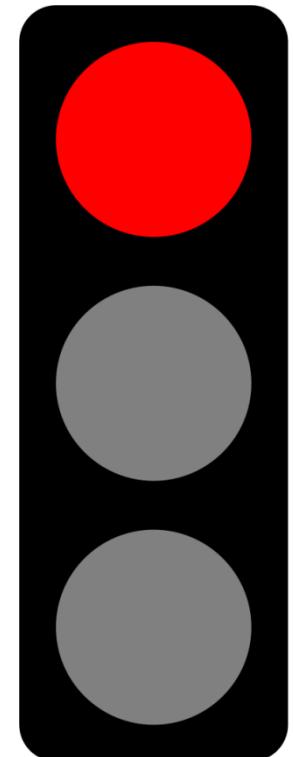
A



B



A



Don't Program lock-free: ABA

A lock-free, singly linked list (stack)



- Thread 1
 - wants to remove A
 - stores
 - head = A
 - next = B
 - checks if A == head
 - make B to the new head
 - B is already deleted by Thread 2
- Thread 2
 - removes A
 - removes B and deletes B
 - pushes A back

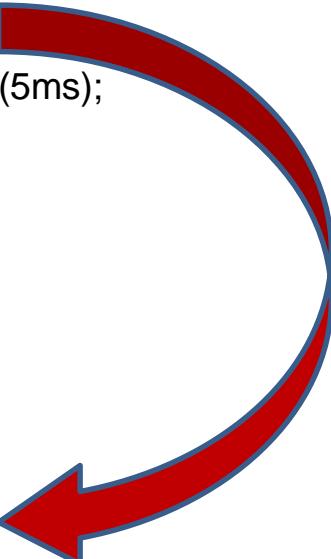
Use Proven Patterns

Wait with sequentiel consistency

```
std::vector<int> mySharedWork;
std::atomic<bool> dataReady(false);

void waitingForWork(){
    while ( !dataReady.load() ){
        std::this_thread::sleep_for(5ms);
    }
    mySharedWork[1] = 2;
}

void setDataReady(){
    mySharedWork = {1, 0, 3};
    dataReady.store(true);
}
```



```
int main(){
    std::thread t1(waitingForWork);
    std::thread t2(setDataReady);
    t1.join();
    t2.join();
    for (auto v: mySharedWork){
        std::cout << v << " ";           // 1 2 3
    }
};
```

Use Proven Patterns

Wait with acquire-release semantic

```
std::vector<int> mySharedWork;
std::atomic<bool> dataReady(false);

void waitForWork(){
    while ( !dataReady.load(std::memory_order_acquire) ){
        std::this_thread::sleep_for(5ms);
    }
    mySharedWork[1] = 2;
}

void setDataReady(){
    mySharedWork = {1, 0, 3};
    dataReady.store(true, std::memory_order_release);
}
```



```
int main(){

    std::thread t1(waitForWork);
    std::thread t2(setDataReady);
    t1.join();
    t2.join();
    for (auto v: mySharedWork){
        std::cout << v << " ";      // 1 2 3
    }
};
```

Use Proven Patterns

Atomic counter

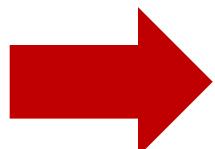
```
#include <vector>
#include <iostream>
#include <thread>
#include<atomic>

std::atomic<int> count{0};

void add() {
    for (int n = 0; n < 1000; ++n) {
        count.fetch_add(1, std::memory_order_relaxed);
    }
}

int main() {
    std::vector<std::thread> v;
    for (int n = 0; n < 10; ++n) {
        v.emplace_back(add);
    }
    for (auto& t : v) { t.join(); }
    std::cout << count;      // 10000
}
```

Don't Reinvent the Wheel



[Boost.Lockfree](#)
[CDS \(Concurrent Data Structures\)](#)

Don't Reinvent the Wheel

- Boost.Lockfree
 - Queue
 - A lock-free multi-producer/multi-consumer queue
 - Stack
 - A lock-free multi-producer/multi-consumer stack
 - spsc_queue
 - A wait-free single-producer/single-consumer queue (commonly known as ringbuffer)

Don't Reinvent the Wheel

- Concurrent Data Structures (CDS)
 - Contains a lot of intrusive and non-intrusive containers
 - Stacks (lock-free)
 - Queues and Priority-Queues (lock-free)
 - Ordered lists
 - Ordered sets and maps (lock-free and lock-based)
 - Unordered sets and maps (lock-free and lock-based)

Blogs

www.grimm-jaud.de [De]

www.ModernesCpp.com [En]

Rainer Grimm
Training, Coaching, and
Technology Consulting
www.ModernesCpp.de