

C++ Concurrency - Formalised

Salomon Sickert

Technische Universität München

26th April 2013

Mutex Algorithms

- At most one thread is in the critical section at any time.

Dekker's Mutex Algorithm [2]

Initialisation

```
1 x = 0; y = 0;
```

Thread 1

```
1 x = 1;  
2 if (y == 1) {  
3   ... //Busy Wait  
4 }  
5 // Critical Section
```

Thread 2

```
1 y = 1;  
2 if (x == 1) {  
3   ... //Busy Wait  
4 }  
5 // Critical section
```

Dekker's Mutex Algorithm [2]

Initialisation

```
1 x = 0; y = 0;
```

Thread 1

```
1 x = 1;  
2 if (y == 1) {  
3   ... //Busy Wait  
4 }  
5 // Critical Section
```

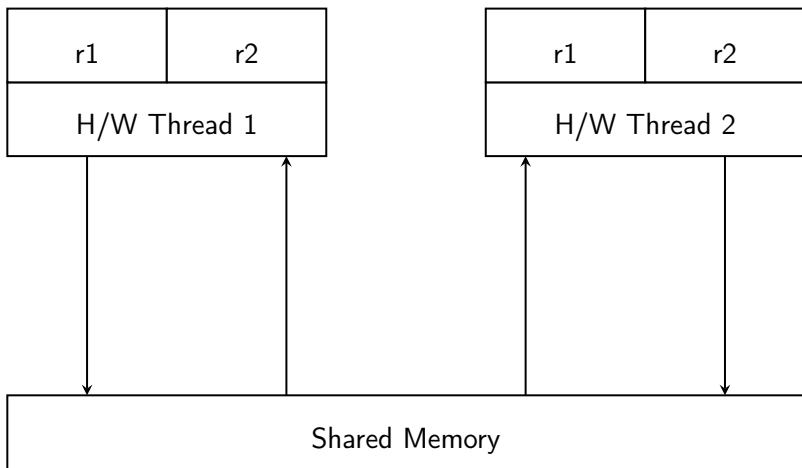
```
1 MOV [x] <- 1  
2 MOV r1 <- [y]
```

Thread 2


```
1 y = 1;  
2 if (x == 1) {  
3   ... //Busy Wait  
4 }  
5 // Critical section
```

```
1 MOV [y] <- 1  
2 MOV r2 <- [x]
```


Modern x86 Multiprocessors - simplified (based on [4])




Modern x86 Multiprocessors - simplified




Modern x86 Multiprocessors and Dekker's Algorithm




Modern x86 Multiprocessors and Dekker's Algorithm



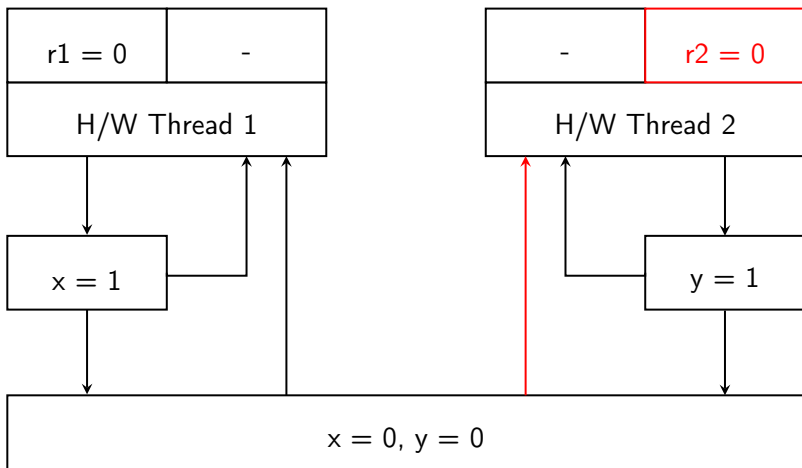
Modern x86 Multiprocessors and Dekker's Algorithm



Modern x86 Multiprocessors and Dekker's Algorithm



Modern x86 Multiprocessors and Dekker's Algorithm



Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**

Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**
- Due to write buffers threads may see a **subtly different memory.**

Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**
- Due to write buffers threads may see a **subtly different memory**.
- The processor exhibits a *relaxed memory model*.

Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**
- Due to write buffers threads may see a **subtly different memory**.
- The processor exhibits a *relaxed memory model*.
- Solutions:
 - Assembler: FENCE instructions.
 - C++11: Special types. (Atomics)

Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**
- Due to write buffers threads may see a **subtly different memory**.
- The processor exhibits a *relaxed memory model*.
- Solutions:
 - Assembler: FENCE instructions.
 - C++11: Special types. (Atomics)
- Sequential consistency is restored by paying a performance penalty.

Modern x86 Multiprocessors

- Both threads may enter the critical section at the **same time!**
- Due to write buffers threads may see a **subtly different memory**.
- The processor exhibits a *relaxed memory model*.
- Solutions:
 - Assembler: FENCE instructions.
 - C++11: Special types. (Atomics)
- Sequential consistency is restored by paying a performance penalty.
- Some Non-x86 architectures exhibit even weaker models, e.g ARM.

Mathematizing C++ Concurrency ([3])

- Given a program p , what are the possible ways to execute it?

Mathematizing C++ Concurrency ([3])

- Given a program p , what are the possible ways to execute it?
 1. Calculate X_{opsem} from p .
 - Loosely speaking: Structure of the program.

Mathematizing C++ Concurrency ([3])

- Given a program p , what are the possible ways to execute it?
 1. Calculate X_{opsem} from p .
 - Loosely speaking: Structure of the program.
 2. Find all $X_{witness}$ consistent with X_{opsem} .
 - Loosely speaking: Different executions of the program.

Mathematizing C++ Concurrency ([3])

- Given a program p , what are the possible ways to execute it?
 1. Calculate X_{opsem} from p .
 - Loosely speaking: Structure of the program.
 2. Find all $X_{witness}$ consistent with X_{opsem} .
 - Loosely speaking: Different executions of the program.
 3. Check for undefined behaviour.
 - Reading from uninitialized variables
 - Unsequenced Races (e.g. $x == (x = 2)$)
 - Data Races
 - ...

X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4    ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```

X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4    ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```

e:W x=1

X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4    ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```

f:W y=2

e:W x=1

X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4    ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```

f:W y=2

e:W x=1

a:W a=1


c:R y=?

b:R x=?

d:W z=(x? == y?)


X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```



X_{opsem} : An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```



- Independent from the architecture


- Independent from the architecture
- Composed of (among other parts):
 - Actions (simplified)
 - $aid : R \mid = v$
 - $aid : W \mid = v$

- Independent from the architecture
- Composed of (among other parts):
 - Actions (simplified)
 - $aid : R \mid = v$
 - $aid : W \mid = v$
 - Binary relations:
 - $\xrightarrow{\text{sequenced-before}} (sb)$
 - $\xrightarrow{\text{additional-synchronized-with}} (asw)$

- Independent from the architecture
- Composed of (among other parts):
 - Actions (simplified)
 - $aid : R \mid = v$
 - $aid : W \mid = v$
 - Binary relations:
 - $\xrightarrow{\text{sequenced-before}} (sb)$
 - $\xrightarrow{\text{additional-synchronized-with}} (asw)$
- In this special case: $\xrightarrow{\text{simple-happens-before}} =$
 $(\xrightarrow{\text{sequenced-before}} \cup \xrightarrow{\text{additional-synchronized-with}})^+$


$X_{witness}$: An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```




$X_{witness}$: An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```



$X_{witness}$: An Example

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```



$X_{witness}$

- Dependent on the architecture

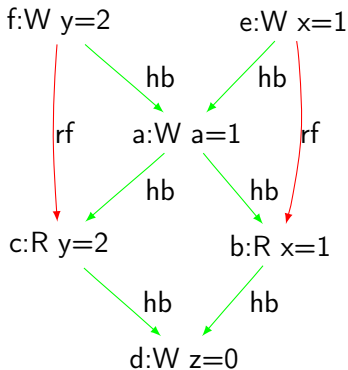
$X_{witness}$

- Dependent on the architecture
- Composed of:
 - Binary relations:
 - $\xrightarrow{\text{reads-from}}$ (rf)

- Dependent on the architecture
- Composed of:
 - Binary relations:
 - $\xrightarrow{\text{reads-from}}$ (rf)
 - $\xrightarrow{\text{sequentialconsistency}}$ (sc) (not applicable in the example)
 - $\xrightarrow{\text{modificationorder}}$ (mo) (not applicable in the example)


$X = (X_{opsem}, X_{witness})$: An execution candidate

```
1 int main() {  
2   int x, y, z, a;  
3   {{{ x = 1;  
4     ||| y = 2;  
5   }}};  
6   a = 1;  
7   z = (x == y);  
8   return 0;  
9 }
```



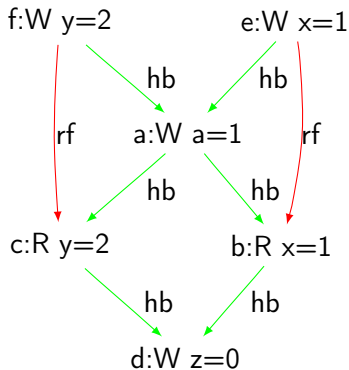
$X = (X_{opsem}, X_{witness})$: Undefined Behaviour?

- Uninitialised Reads?
- Unsequenced Races?
- Data Races?




$X = (X_{opsem}, X_{witness})$: Undefined Behaviour?

- Uninitialised Reads? \times
- Unsequenced Races?
- Data Races?




$X = (X_{opsem}, X_{witness})$: Undefined Behaviour?

- Uninitialised Reads? ×
- Unsequenced Races? ×
- Data Races?




$X = (X_{opsem}, X_{witness})$: Undefined Behaviour?

- Uninitialised Reads? ×
- Unsequenced Races? ×
- Data Races? ×




Undefined Behaviour: Data Races

```
1 int main() {  
2   int x = 0;  
3   {{{ x = 1;  
4     ||| x = 2;  
5   }}};  
6   return x;  
7 }
```




Undefined Behaviour: Data Races

```
1 int main() {  
2   int x = 0;  
3   {{{ x = 1;  
4     ||| x = 2;  
5   }}};  
6   return x;  
7 }
```



Undefined Behaviour: Data Races

```
1 int main() {  
2   int x = 0;  
3   {{{ x = 1;  
4     ||| x = 2;  
5   }}};  
6   return x;  
7 }
```



The Formalised C++ Memory Model

```
cpp_memory_model opsem (p : program) =  
  let pre_executions = {(Xopsem, Xwitness) |  
    opsem p Xopsem ∧ consistent_execution Xopsem Xwitness}} in  
  if ∃ X ∈ pre_executions. undefined_behaviour X  
  then NONE  
  else SOME pre_executions
```

C++11 Concurrency / Language Features

Concurrency Idioms (**Atomics**, Mutexes, Threads)

C++11 Concurrency / Language Features

Concurrency Idioms (**Atomics**, Mutexes, Threads)

Pre-C++11

- No memory model for multi-threaded code
- Concurrency Idioms provided by a third party:
 - pthreads
 - OpenMP
- Drawbacks: No formalised standard, compiler may produce incorrect code.

C++11 Concurrency / Language Features

Concurrency Idioms (**Atomics**, Mutexes, Threads)

Pre-C++11

- No memory model for multi-threaded code
- Concurrency Idioms provided by a third party:
 - pthreads
 - OpenMP
- Drawbacks: No formalised standard, compiler may produce incorrect code.

C++11

- A memory model for multi-threaded code
- Concurrency Idioms in are part of the language (`std::atomic<T>` [1], `std::mutex`, `std::thread`)
 - Similar to Java
- Benefits: Compiler is able to produce correct code.

Applications of the Formal Memory Model

- Corrections to the C++0x standard.
 - `memory_order_seq_cst` was in fact not sequentially consistent

Applications of the Formal Memory Model

- Corrections to the C++0x standard.
 - `memory_order_seq_cst` was in fact not sequentially consistent
- Confidence in memory model and specification.

Applications of the Formal Memory Model

- Corrections to the C++0x standard.
 - `memory_order_seq_cst` was in fact not sequentially consistent
- Confidence in memory model and specification.
- Verify correctness of prototype implementations.




Figure: Figure taken from [3]

Applications of the Formal Memory Model

- Corrections to the C++0x standard.
 - `memory_order_seq_cst` was in fact not sequentially consistent
- Confidence in memory model and specification.
- Verify correctness of prototype implementations.

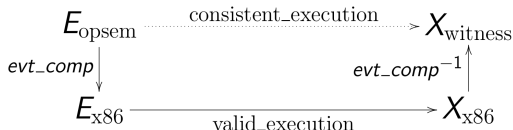
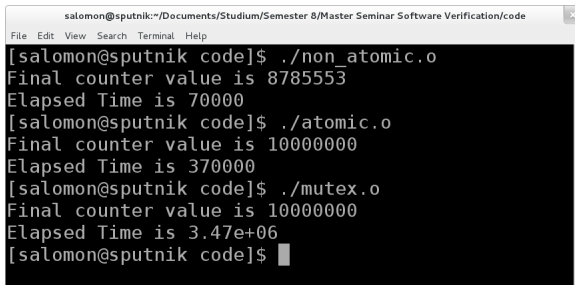


Figure: Figure taken from [3]

- Developer support.

Questions?

Bonus: Atomics vs. Mutexes - short presentation







```
salomon@sputnik:~/Documents/Studium/Semester 8/Master Seminar Software Verification/code
File Edit View Search Terminal Help
[salomon@sputnik code]$ ./non_atomic.o
Final counter value is 8785553
Elapsed Time is 70000
[salomon@sputnik code]$ ./atomic.o
Final counter value is 10000000
Elapsed Time is 370000
[salomon@sputnik code]$ ./mutex.o
Final counter value is 10000000
Elapsed Time is 3.47e+06
[salomon@sputnik code]$
```

Bonus: Dekker's algorithm in CppMem

`http://svr-pes20-cppmem.cl.cam.ac.uk/cppmem/index.html`

References

-  C++ atomic types and operations, <http://www.open-std.org/jtc1/sc22/wg21/docs/papers/2007/n2427.html>.
-  Dekker's algorithm, <http://www.cs.utexas.edu/users/ewd/transcriptions/ewd01xx/ewd123>.
-  Mark Batty, Scott Owens, Susmit Sarkar, Peter Sewell, and Tjark Weber.
Mathematizing c++ concurrency.
SIGPLAN Not., 46(1):55–66, January 2011.
-  Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and Magnus O. Myreen.
x86-tso: a rigorous and usable programmer's model for x86 multiprocessors.
Commun. ACM, 53(7):89–97, July 2010.