T. Hoefler, M. Pueschel

Lecture 3: Memory Models

Teaching assistant: Salvatore Di Girolamo

Motivational video: https://www.youtube.com/watch?v=tW2hT0g4OUs

INTERESTING CASE OF RULE 2

Based on the presented data, one may conclude that using **-O3** is always a good idea.

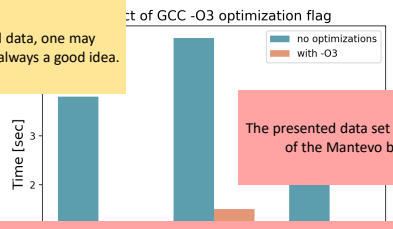

The presented data set contains only **a subset** of the Mantevo benchmark suite.

The incompleteness of data may lead to wrong conclusions. Sometimes **-O3** may not be a good idea for a code: e.g., vectorization (enabled by **-O3**) may segfault on a loop which does unaligned memory access on x86. But this is not demonstrated by the presented dataset.

Review of last lecture

- **Architecture case studies**
  - Memory performance is often the bottleneck
  - Parallelism grows with compute performance
  - Caching is important
  - Several issues to address for parallel systems
- **Cache Coherence**
  - Hardware support to aid programmers
  - Two guarantees:
    Write propagation (updates are eventually visible to all readers)
    Write serialization (writes to the same location are observed in global order)
  - Two major mechanisms:
    Snooping
    Directory-based
  - Protocols: MESI (MOESI, MESIF)

DPHPC Overview

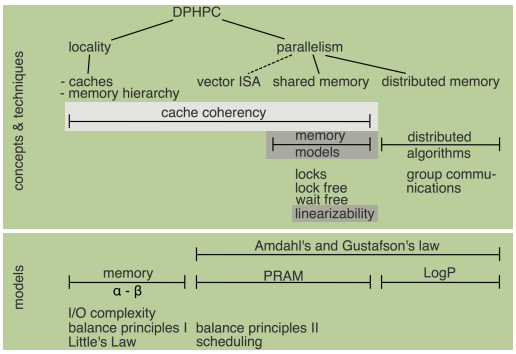

Goals of this lecture

- **Don't forget the projects!**
  - Groups to be defined by Thursday (send email to Salvatore)
  - Project progress presentations on 11/6 (<1 month from now)!
- **Cache-coherence is not enough!**
  - Many more subtle issues for parallel programs!
- **Memory Models**
  - Sequential consistency
  - Why threads cannot be implemented as a library ☺
  - Relaxed consistency models
- **Linearizability**
  - More complex objects

Is coherence everything?

- **Coherence is concerned with behavior of *individual* locations**
- **Consider the program (initial X,Y,Z = 0)**

```
P1          P2
Y=10        while (X==0);
X=2         Z=Y
```

  - Class question: what value will Z on P2 have?
- **Y=10 does not need to have completed before X=2 is visible to P2!**
  - This allows P2 to exit the loop and read Y=0
  - This may not be the intent of the programmer!
  - This may be due to congestion (imagine X is pushed to a remote cache while Y misses to main memory) and or due to write buffering, or …
- **Bonus class question: what happens when Y and X are on the same cache line (assume simple MESI and no write buffer)?**

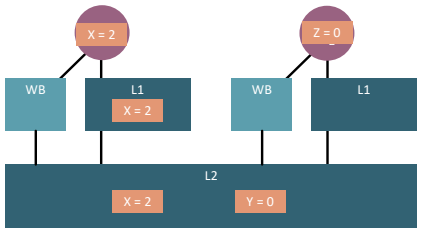

Memory Models

- **Need to define what it means to "read a location" and "to write a location" and the respective ordering!**
  - What values should be seen by a processor
- **First thought: extend the abstractions seen by a sequential processor:**
  - **Compiler** and **hardware** maintain data and control dependencies at all levels:

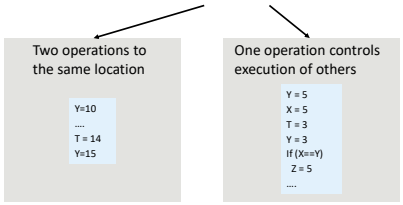

Sequential Processor

- **Correctness condition:**
  - The result of the execution is the same as if the operations had been executed in the order specified by the program
    "program order"
  - A read returns the value last written to the same location
    "last" is determined by program order!
- **Consider only memory operations (e.g., a trace)**
- **N Processors**
  - P1, P2, …., PN
- **Operations**
  - Read, Write on shared variables (initial state: most often all 0)
- **Notation:**
  - P1: R(x):3 P1 reads x and observes the value 3
  - P2: W(x,5) P2 writes 5 to variable x

Terminology

- **Program order**
  - Deals with a *single* processor
  - Per-processor order of memory accesses, determined by program's *Control flow*
  - Often represented as trace
- **Visibility order**
  - Deals with operations on *all* processors
  - Order of memory accesses observed by one or more processors
  - E.g., "every read of a memory location returns the value that was written last"
    Defined by memory model

Memory Models

- **Contract at each level between programmer and processor**

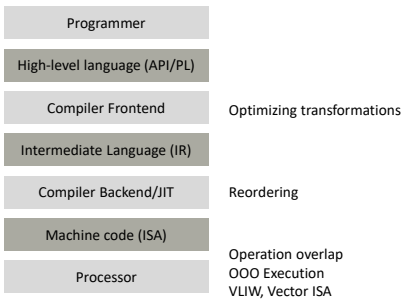

Sequential Consistency

- **Extension of sequential processor model**
- **The execution happens as if**
  - The operations of all processes were executed in some sequential order (atomicity requirement), and
  - The operations of each individual processor appear in this sequence in the order specified by the program (program order requirement)
- **Applies to all layers!**
  - Disallows many compiler optimizations (e.g., reordering of *any* memory instruction)
  - Disallows many hardware optimizations (e.g., store buffers, nonblocking reads, invalidation buffers)

Illustration of Sequential Consistency

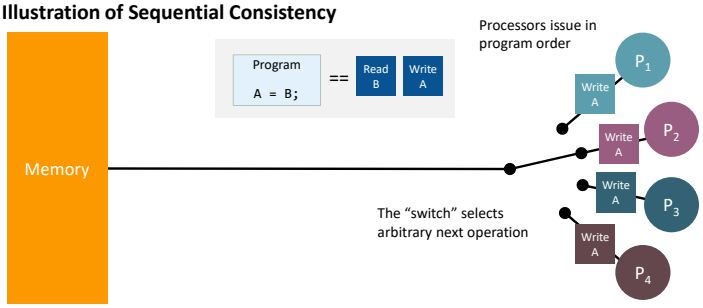

- **Globally consistent view of memory operations (atomicity)**
- **Strict ordering in program order**

Original SC Definition

"The result of any execution is the same as if the operations of all the processes were executed in some sequential order and the operations of each individual process appear in this sequence in the order specified by its program"

(Lamport, 1979)

Good read: Gharachorloo et al.: "Memory consistency and event ordering in scalable shared-memory multiprocessors."

Alternative SC Definition

- **Textbook: Hennessy/Patterson Computer Architecture**
- **A sequentially consistent system maintains three invariants:**
  1. A load L from memory location A issued by processor P_i obtains the value of the previous store to A by P_i, unless another processor has to stored a value to A in between
  2. A load L from memory location A obtains the value of a store S to A by another processor P_k if S and L are "sufficiently separated in time" and if no other store occurred between S and L
  3. Stores to the same location are serialized (defined as in (2))
- **"Sufficiently separated in time" not precise**
  - Works but is not formal (a formalization must include all possibilities)

Example Operation Reordering

- **Recap: "normal" sequential assumption:**
  - Compiler and hardware can reorder instructions as long as control and data dependencies are met
- **Examples:**

Compiler
- Register allocation
- Code motion
- Common subexpression elimination
- Loop transformations

Hardware
- Pipelining
- Multiple issue (OOO)
- Write buffer bypassing
- Nonblocking reads

Simple compiler optimization

- **Initially, all values are zero**

P1	P2
input = 23 ready = 1	while (ready == 0) {} compute(input)

- Assume P1 and P2 are compiled separately!
- What optimizations can a compiler perform for P1?
  Register allocation or even replace with constant, or Switch statements
- What happens?
  P2 may never terminate, or Compute with wrong input

Sequential Consistency Examples

- **Relying on program order: Dekker's algorithm**
  - Initially, all zero

P1	P2

a = 1
if(b == 0)
  critical section
  a = 0
	b = 1
	if(a == 0)
	  critical section
	  b = 0

- What can happen at compiler and hardware level?

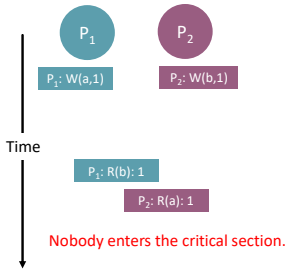

Nobody enters the critical section.

Alternatively, both writes may have went to a write buffer, in which case both Ps would read 0 and enter the critical section together.

Sequential Consistency Examples

- **Relying on single sequential order (atomicity): three sharers**

P1	P2	P3

a = 5
a = 1
	if (a == 1)
	  b = 1
		if (b == 1)
		  print(a)

- What can be printed if visibility is not atomic?

What each **P** thinks the order is:

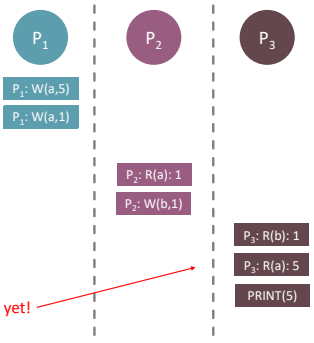

P3 has not seen P1: W(a,1) yet!

Optimizations violating program order

- **Analyzing P1 and P2 in isolation!**
  - Compiler can reorder

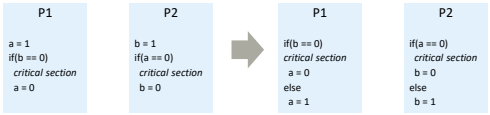

  - Hardware can reorder, assume writes of a,b go to write buffer or speculation

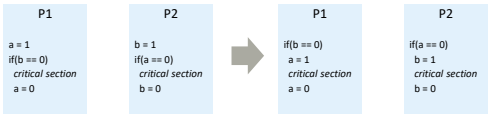

Considerations

- **Define partial order on memory requests A → B**
  - If P_i issues two requests A and B and A is issued before B in program order, then A → B
  - A and B are issued to the same variable, and A is issued first, then A → B (on all processors)
- **These partial orders can be interleaved, define a total order**
  - Many total orders are sequentially consistent!
- **Example:**
  - P1: W(a), R(b), W(c)
  - P2: R(a), W(a), R(b)
  - Are the following schedules (total orders) sequentially consistent?
    1. *P1:W(a), P2:R(a), P2:W(a), P1:R(b), P2:R(b), P1:W(c)*
    2. *P1:W(a), P2:R(a), P1:R(b), P2:R(b), P1:W(c), P2:W(a)*
    3. *P2:R(a), P2:W(a), P1:R(b), P1:W(a), P1:W(c), P2:R(b)*

Write buffer example

- **Write buffer**
  - Absorbs writes faster than the next cache → prevents stalls
  - Aggregates writes to the same cache line → reduces cache traffic

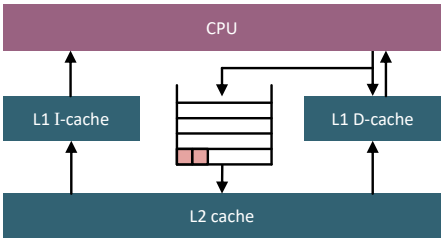

Write buffer example

- **Reads can bypass previous writes for faster completion**
  - If read and write access different locations
  - No order between write and following read (W → R)

P1	P2
W(a,1)	W(b,1)
R(b):0	R(a):0

not seq. consistent

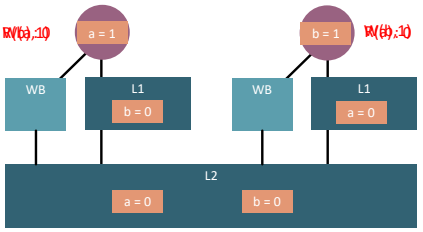

Nonblocking read example

- **W → W: OK**
- **R → W, R → R: No order between read and following read/write**

P1	P2
R(y):2	W(x,1)
R(x):0	W(y,2)

not seq. consistent

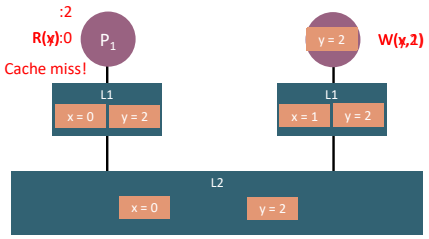

Discussion

- **Programmer's view:**
  - Prefer sequential consistency
  - Easiest to reason about
- **Compiler/hardware designer's view:**
  - Sequential consistency disallows many optimizations!
  - Substantial speed difference
  - ➢ Most architectures and compilers don't adhere to sequential consistency!
- **Solution: synchronized programming**
  - Access to shared data (aka. "racing accesses") are ordered by synchronization operations
  - Synchronization operations guarantee memory ordering (aka. fence)
  - More later!

Cache Coherence vs. Memory Model

- **Varying definitions!**

- **Cache coherence: a mechanism that propagates writes to other processors/caches if needed, recap:**
  - Writes are eventually visible to all processors
  - Writes to the same location are observed in (one) order

- **Memory models: define the bounds on when the value is propagated to other processors**
  - E.g., sequential consistency requires *all* reads and writes to be ordered in program order

Good read: McKenney: "Memory Barriers: a Hardware View for Software Hackers"

The fun begins: Relaxed Memory Models

- **Sequential consistency**
  - R→R, R→W, W→R, W→W (all orders guaranteed)

- **Relaxed consistency (varying terminology):**
  - Processor consistency (aka. TSO)
    Relaxes W→R
  - Partial write (store) order (aka. PSO)
    Relaxes W→R, W→W
  - Weak consistency and release consistency (aka. RMO)
    Relaxes R→R, R→W, W→R, W→W
  - Other combinations/variants possible
    There are even more types of orders (above is a simplification)

Architectures

Memory ordering in some architectures[2][3]

Type	Alpha	ARMv7	PA-RISC	POWER	SPARC RMO	SPARC PSO	SPARC TSO	x86	x86 oostore	AMD64	IA-64	zSeries
Loads reordered after loads	Y	Y	Y	Y	Y				Y		Y	
Loads reordered after stores	Y	Y	Y	Y	Y				Y		Y	
Stores reordered after stores	Y	Y	Y	Y	Y	Y			Y		Y	
Stores reordered after loads	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y	Y
Atomic reordered with loads	Y	Y		Y	Y						Y	
Atomic reordered with stores	Y	Y		Y	Y	Y					Y	
Dependent loads reordered	Y											
Incoherent Instruction cache pipeline	Y	Y		Y	Y	Y	Y	Y	Y		Y	Y

Some older x86 and AMD systems have weaker memory ordering[4]

Source: wikipedia

Case Study: Memory ordering on Intel (x86)

- **Intel® 64 and IA-32 Architectures Software Developer's Manual**
  - Volume 3A: System Programming Guide
  - Chapter 8.2 Memory Ordering
  - http://www.intel.com/products/processor/manuals/

- **Google Tech Talk: IA Memory Ordering**
  - Richard L. Hudson
    http://www.youtube.com/watch?v=WUfvvFD5tAA

x86 Memory model: TLO + CC

- **Total lock order (TLO)**
  - Instructions with "lock" prefix enforce total order across all processors
  - Implicit locking: xchg (locked compare and exchange)

- **Causal consistency (CC)**
  - Write visibility is transitive

- **Eight principles**
  - After some revisions ☺

The Eight x86 Principles

1. **"Reads are not reordered with other reads." (R→R)**
2. **"Writes are not reordered with other writes." (W→W)**
3. **"Writes are not reordered with older reads." (R→W)**
4. **"Reads may be reordered with older writes to different locations but not with older writes to the same location." (NO W→R!)**
5. **"In a multiprocessor system, memory ordering obeys causality." (memory ordering respects transitive visibility)**
6. **"In a multiprocessor system, writes to the same location have a total order." (implied by cache coherence)**
7. **"In a multiprocessor system, locked instructions have a total order." (enables synchronized programming!)**
8. **"Reads and writes are not reordered with locked instructions. " (enables synchronized programming!)**

Principle 1 and 2

Reads are not reordered with other reads. (R→R)
Writes are not reordered with other writes. (W→W)

All values zero initially. r1 and r2 are registers.

P1	P2
a = 1 b = 2	r1 = b r2 = a

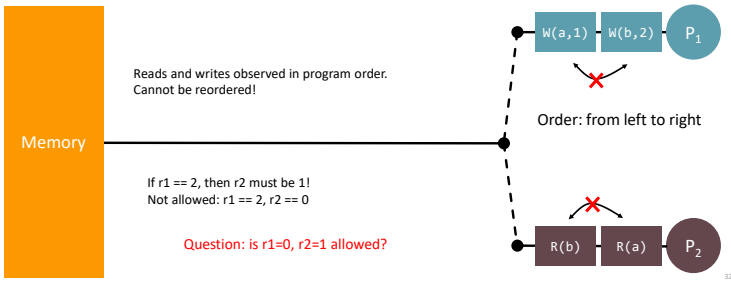

Principle 3

Writes are not reordered with older reads. (R→W)

P1	P2
r1 = a b = 1	r2 = b a = 1

All values zero initially

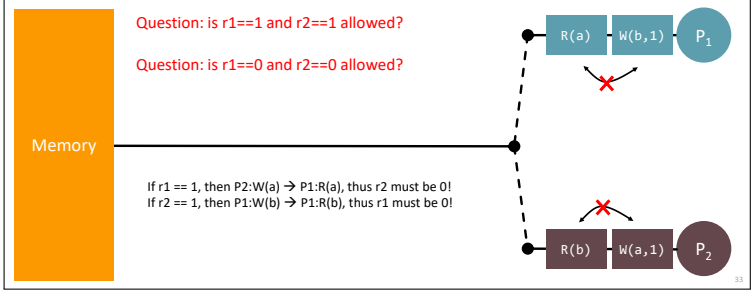

Principle 4

Reads may be reordered with older writes to different locations but not with older writes to the same location. (NO W→R!)

All values zero initially

P1	P2
a = 1 r1 = b	b = 1 r2 = a

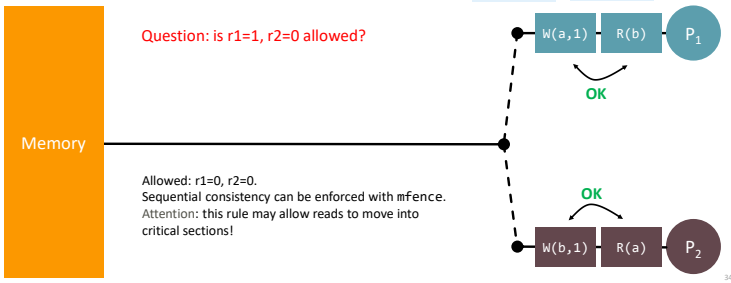

Principle 5

In a multiprocessor system, memory ordering obeys causality (memory ordering respects transitive visibility).

All values zero initially

P1	P2	P3
a = 1	r1 = a b = 1	r2 = b r3 = a

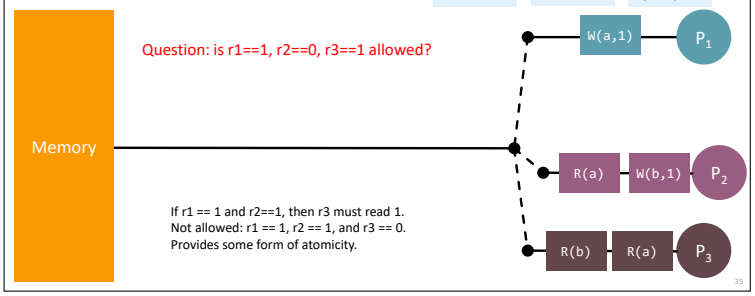

Principle 6

In a multiprocessor system, writes to the same location have a total order (implied by cache coherence).

All values zero initially

P1	P2	P3	P4
a=1	a=2	r1 = a r2 = a	r3 = a r4 = a

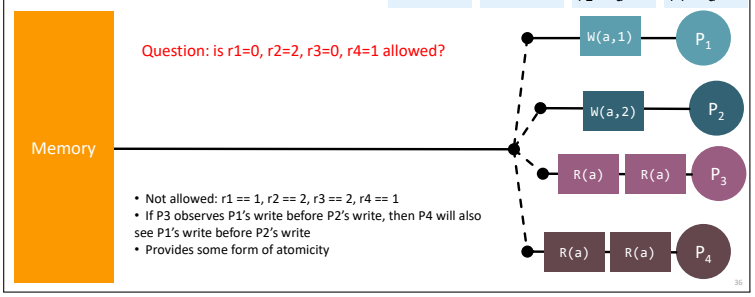

Principle 7

In a multiprocessor system, locked instructions have a total order. (enables synchronized programming!)

All values zero initially, registers r1==r2==1

P1	P2	P3	P4
xchg(a,r1)	xchg(b,r2)	r3 = a r4 = b	r5 = b r6 = a

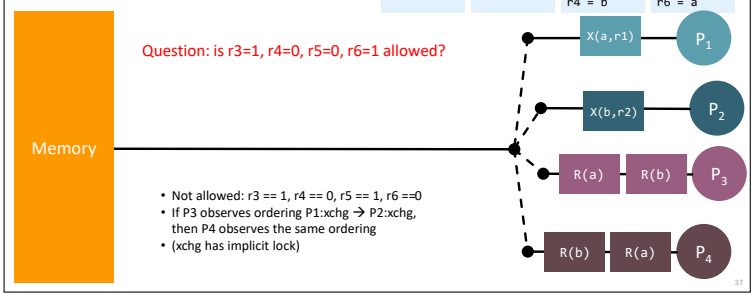

Principle 8

Reads and writes are not reordered with locked instructions. (enables synchronized programming!)

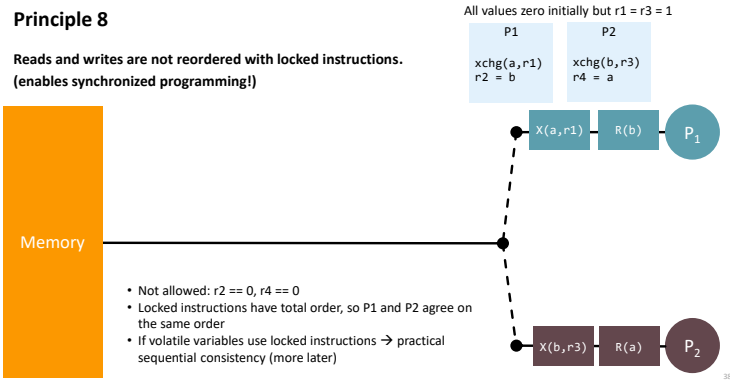

- Not allowed: r2 == 0, r4 == 0
- Locked instructions have total order, so P1 and P2 agree on the same order
- If volatile variables use locked instructions → practical sequential consistency (more later)

An Alternative View: x86-TSO

- **Sewell el al.: "x86-TSO: A Rigorous and Usable Programmer's Model for x86 Multiprocessors", CACM May 2010**

"[…] **real multiprocessors typically do not provide the sequentially consistent memory** that is assumed by most work on semantics and verification. Instead, they have relaxed memory models, varying in subtle ways between processor families, in which different hardware threads may have only loosely consistent views of a shared memory. Second, **the public vendor architectures, supposedly specifying what programmers can rely on, are often in ambiguous informal prose (a particularly poor medium for loose specifications), leading to widespread confusion.** […] We present a new x86-TSO programmer's model that, to the best of our knowledge, suffers from none of these problems. **It is mathematically precise** (rigorously defined in HOL4) but can be presented as an **intuitive abstract machine which should be widely accessible to working programmers.** […]"

Newer RMA systems: A. Dan, P. Lam, TH, A. Vechev: Modeling and Analysis of Remote Memory Access Programming, ACM OOPSLA'16

Notions of Correctness

- **We discussed so far:**
  - Read/write of the same location
    Cache coherence (write serialization and atomicity)
  - Read/write of multiple locations
    Memory models (visibility order of updates by cores)

- **Now: objects (variables/fields with invariants defined on them)**
  - Invariants "tie" variables together
  - Sequential objects
  - Concurrent objects

Sequential Objects

- **Each object has a type**
- **A type is defined by a class**
  - Set of fields forms the state of an object
  - Set of methods (or free functions) to manipulate the state

- **Remark**
  - An Interface is an abstract type that defines behavior
    A class implementing an interface defines several types

Running Example: FIFO Queue

- **Insert elements at tail**
- **Remove elements from head**
  - Initial: `head = tail = 0`
  - `enq(x)`
  - `enq(y)`
  - `deq() [x]`
  - …

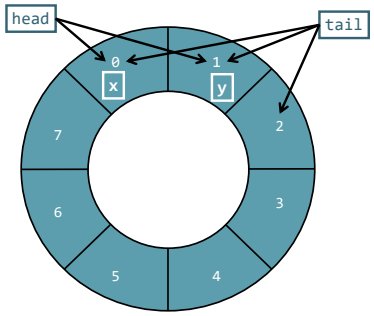

capacity = 8

Sequential Queue

```
class Queue {
private:
  int head, tail;
  std::vector<Item> items;

public:
  Queue(int capacity) {
    head = tail = 0;
    items.resize(capacity);
  }

  // ...

};
```

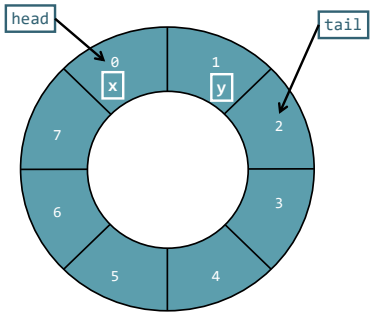

capacity = 8

Sequential Queue

```
class Queue {
  // ...

public:
  void enq(Item x) {
    if((tail+1)%items.size() == head) {
      throw FullException;
    }
    items[tail] = x;
    tail = (tail+1)%items.size();
  }

  Item deq() {
    if(tail == head) {
      throw EmtpyException;
    }
    Item item = items[head];
    head = (head+1)%items.size();
    return item;
  }
};
```

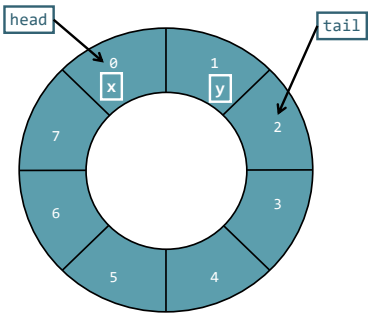

capacity = 8

Sequential Execution

- **(The) one process executes operations one at a time**
  - Sequential ☺
- **Semantics of operation defined by specification of the class**
  - Preconditions and postconditions

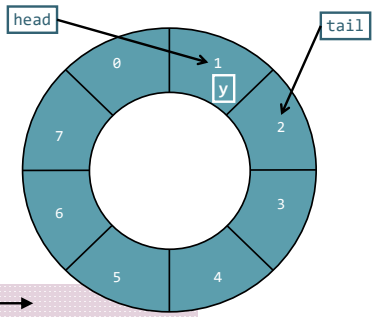

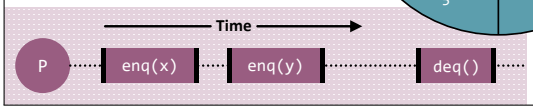

Design by Contract!

- **Preconditions:**
  - Specify conditions that must hold before method executes
  - Involve state and arguments passed
  - Specify obligations a client must meet before calling a method
- **Example: enq()**
  - Queue must not be full!

```
class Queue {
  // ...
  void enq(Item x) {
    assert((tail+1)%items.size() != head);
    // ...
  }
};
```

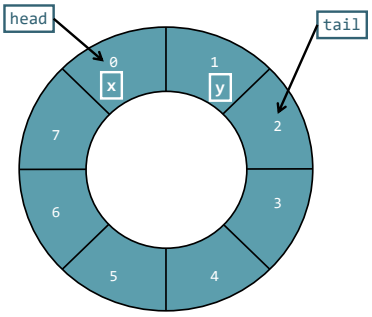

capacity = 8

Design by Contract!

- **Postconditions:**
  - Specify conditions that must hold after method executed
  - Involve old state and arguments passed
- **Example: enq()**
  - Queue must contain element!

```
class Queue {
  // ...
  void enq(Item x) {
    // ...
    assert(
      (tail == (old_tail + 1)%items.size()) &&
      (items[old_tail] == x) );
  }
};
```

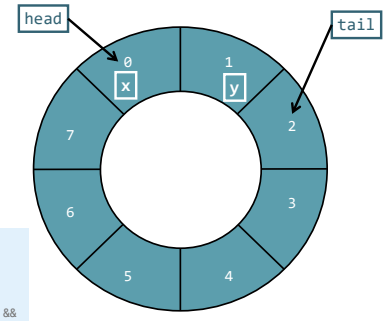

capacity = 8

Sequential specification

- **if(precondition)**
  - Object is in a **specified state**
- **then(postcondition)**
  - The method returns a particular value or
  - Throws a particular exception **and**
  - Leaves the object in a **specified state**
- **Invariants**
  - Specified conditions (e.g., object state) must hold **anytime** a client could invoke an objects method!

Advantages of sequential specification

- **State between method calls is defined**
  - Enables reasoning about objects
  - Interactions between methods captured by side effects on object state
- **Enables reasoning about each method in isolation**
  - Contracts for each method
  - Local state changes global state
- **Adding new methods**
  - Only reason about state changes that the new method causes
  - If invariants are kept: **no need to check old methods**
  - **Modularity!**

Concurrent execution - State

- **Concurrent threads invoke methods on possibly shared objects**
  - At overlapping time intervals!

Property	Sequential	Concurrent
State	Meaningful only between method executions	Overlapping method executions → object may never be "between method executions"

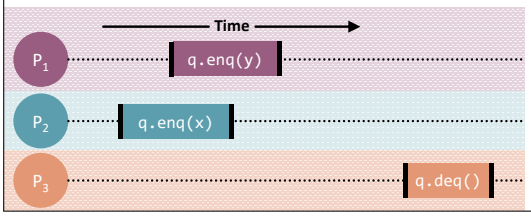

Each method execution takes some non-zero amount of time!

Concurrent execution - Reasoning

- **Reasoning must now include all possible interleavings**
  - Of changes caused by methods themselves

Property	Sequential	Concurrent
Reasoning	Consider each method in isolation; invariants on state before/after execution.	Need to consider all possible interactions; all intermediate states during execution

That is, now we have to consider what will happen if we execute:

- **enq()** concurrently with **enq()**
- **deq()** concurrently with **deq()**
- **deq()** concurrently with **enq()**

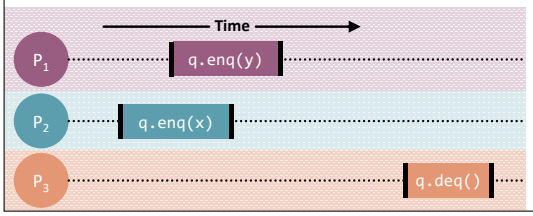

Each method execution takes some non-zero amount of time!

Concurrent execution - Method addition

- **Reasoning must now include all possible interleavings**
  - Of changes caused by and methods themselves

Property	Sequential	Concurrent
Add Method	Without affecting other methods; invariants on state before/after execution.	Everything can potentially interact with everything else

- **Consider adding a method that returns the last item enqueued**

```
Item peek() {
  if(tail == head) throw EmptyException;
  return items[(tail-1) % items.size()];
}
```

```
void enq(Item x) {
  items[tail] = x;
  tail = (tail+1) % items.size();
}
```

```
Item deq() {
  Item item = items[head];
  head = (head+1) % items.size();
}
```

- If **peek()** and **enq()** run concurrently: what if tail has not yet been incremented?
- If **peek()** and **deq()** run concurrently: what if last item is being dequeued?

Concurrent objects

- **How do we describe one?**
  - No pre-/postconditions ☹
- **How do we implement one?**
  - Plan for quadratic or exponential number of interactions and states
- **How do we tell if an object is correct?**
  - Analyze all quadratic or exponential interactions and states

Is it time to panic for (parallel) software engineers?
Who has a solution?

Lock-based queue

```
class Queue {
private:
  int head, tail;
  std::vector<Item> items;
  std::mutex lock;

public:
  Queue(int capacity) {
    head = tail = 0;
    items.resize(capacity);
  }
  // ...
};
```

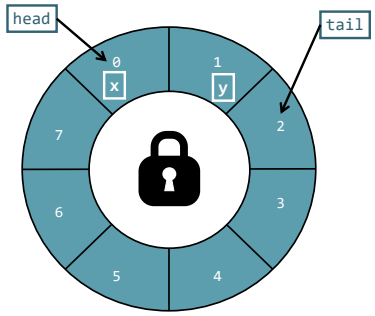

We can use the lock to protect Queue's fields.

Lock-based queue

```
class Queue {
  // ...
public:
  void enq(Item x) {
    std::lock_guard<std::mutex> l(lock);
    if((tail+1)%items.size()==head) {
      throw FullException;
    }
    items[tail] = x;
    tail = (tail+1)%items.size();
  }

  Item deq() {
    std::lock_guard<std::mutex> l(lock);
    if(tail == head) {
      throw EmptyException;
    }
    Item item = items[head];
    head = (head+1)%items.size();
    return item;
  }
};
```

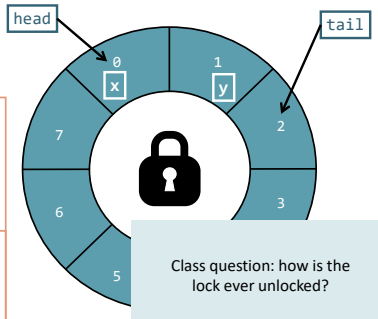

Class question: how is the lock ever unlocked?

One of C++'s ways of implementing a **critical section**

C++ Resource Acquisition is Initialization

- **RAII – suboptimal name**
- **Can be used for locks (or any other resource acquisition)**
  - Constructor grabs resource
  - Destructor frees resource
- **Behaves as if**
  - Implicit unlock at end of block!
- **Main advantages**
  - Always unlock/free lock at exit
  - No “lost” locks due to exceptions or strange control flow (goto ☺)
  - Very easy to use

```
template <typename mutex_impl>
class lock_guard {
  mutex_impl& _mtx; // ref to the mutex

public:
  lock_guard(mutex_impl& mtx ) : _mtx(mtx) {
    _mtx.lock(); // lock mutex in constructor
  }

  ~lock_guard() {
    _mtx.unlock(); // unlock mutex in destructor
  }
};
```


Example execution

... and proceeds.

```
void enq(Item x) {
  std::lock_guard<std::mutex> l(lock);
  if((tail+1)%items.size()==head) {
    throw FullException;
  }
  items[tail] = x;
  tail = (tail+1)%items.size();
}
```

```
Item deq() {
  std::lock_guard<std::mutex> l(lock);
  if(tail == head) {
    throw EmptyException;
  }
  Item item = items[head];
  head = (head+1)%items.size();
  return item;
}
```

enq(x)

deq()

Methods effectively execute one after another, sequentially.

Correctness

- **Is the locked queue correct?**
  - Yes, only one thread has access if locked correctly
  - Allows us again to reason about pre- and postconditions
  - Smells a bit like sequential consistency, no?
- **Class question: What is the problem with this approach?**
  - Same as for SC ☺

It does not scale!
What is the solution here?