

# Alternative Synchronization Strategies — Transactional Memory —

lecture 25 (2021-06-14)

**Master in Computer Science and Engineering**

— Concurrency and Parallelism / 2020-21 —

João Lourenço <joao.lourenco@fct.unl.pt>

# Alternative Synchronization Strategies

- Contents:
  - Coarse-Grained Synchronization
  - Fine-Grained Synchronization
  - Optimistic Synchronization
  - Lazy Synchronization
  - Lock-Free Synchronization
  - Transactional Memory
- Reading list:
  - Chapter 10 of the Textbook
  - Chapter 18 of “The Art of Multiprocessor Programming” by Maurice Herlihy & Nir Shavit (*available at clip*)

# Parallel

# Computing

is here

to stay!

And locks are  
just **not** good  
enough!

# Why locking doesn't scale?

---

- **Not Robust**

- What happens if the thread holding a lock dies?

- Relies on conventions

- Hard to Use

- Conservative

- Deadlocks

- Lost wake-ups

- Not Composable

# Why locking doesn't scale?

- Not Robust
- **Relies on conventions**
  - Lock bit and object bits
  - Exists only in programmer's mind


Actual comment  
from Linux Kernel  
(hat tip: Bradley Kuszmaul)

- Hard to  
  - Cons
  - Dead
  - Lost v
- Not Composable

```
/*  
 * When a locked buffer is visible to the I/O layer  
 * BH_Laundry is set. This means before unlocking  
 * we must clear BH_Laundry,mb() on alpha and then  
 * clear BH_Lock, so no reader can see BH_Laundry set  
 * on an unlocked buffer and then risk to deadlock.  
 */
```

# Why locking doesn't scale?

- Not Robust
- Relies on conventions
- **Hard to use**
  - Conservative
  - Deadlocks
  - Lost wake-ups
- Not composable



# Why locking doesn't scale?

- Not Robust
- Relies on conventions
- Hard to use
  - Conservative
  - Deadlocks
  - Lost wake-ups


```
class Queue {  
  
    /* private fields... */  
    synchronized bool is_empty();  
    synchronized bool is_full();  
    synchronized bool is_enqueue();  
    synchronized bool is_dequeue();  
}
```

## • LOCKS ARE NOT COMPOSABLE!

```
class QueueOperations {  
    synchronized void q_transfer(...);  
}
```

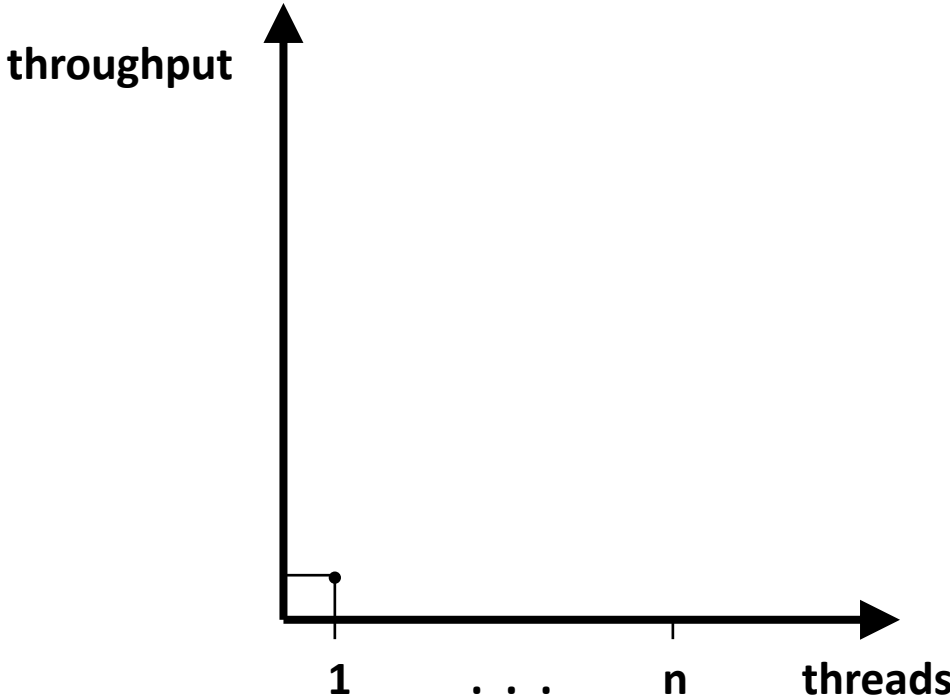
**WRONG!**

# Parallel Throughput

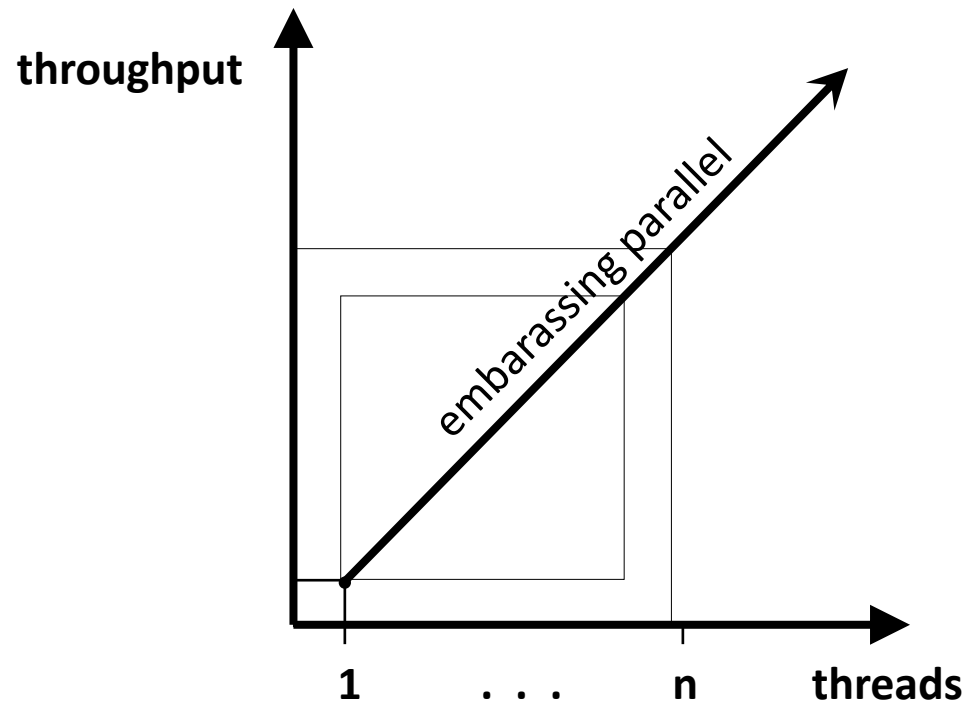




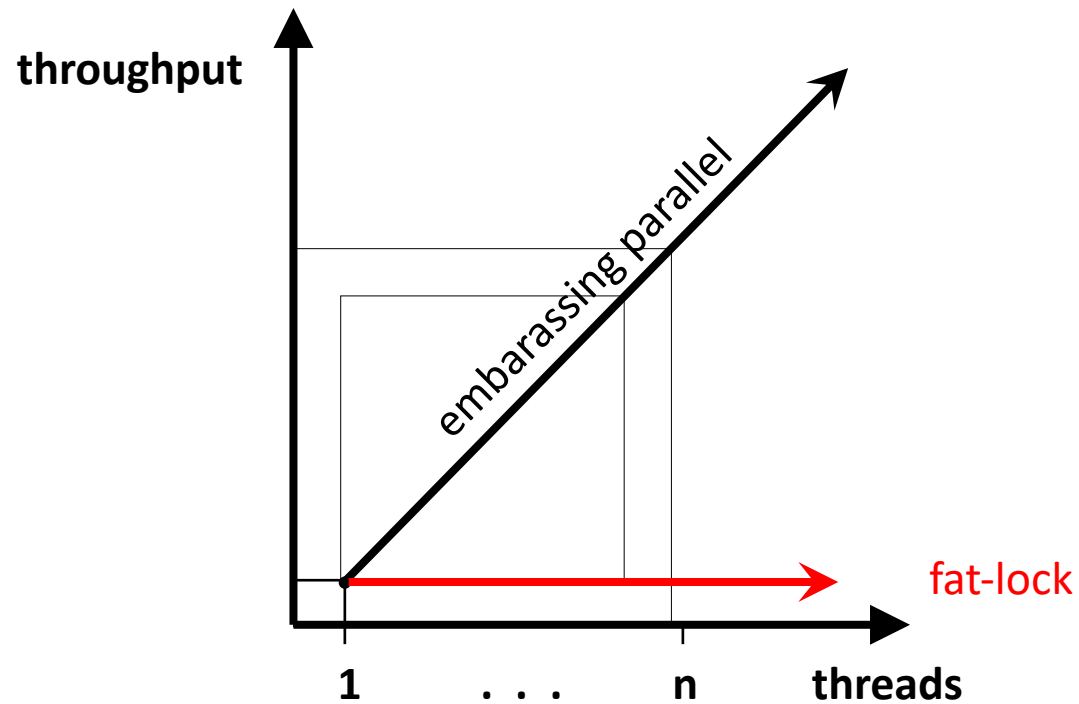
# Parallel Throughput



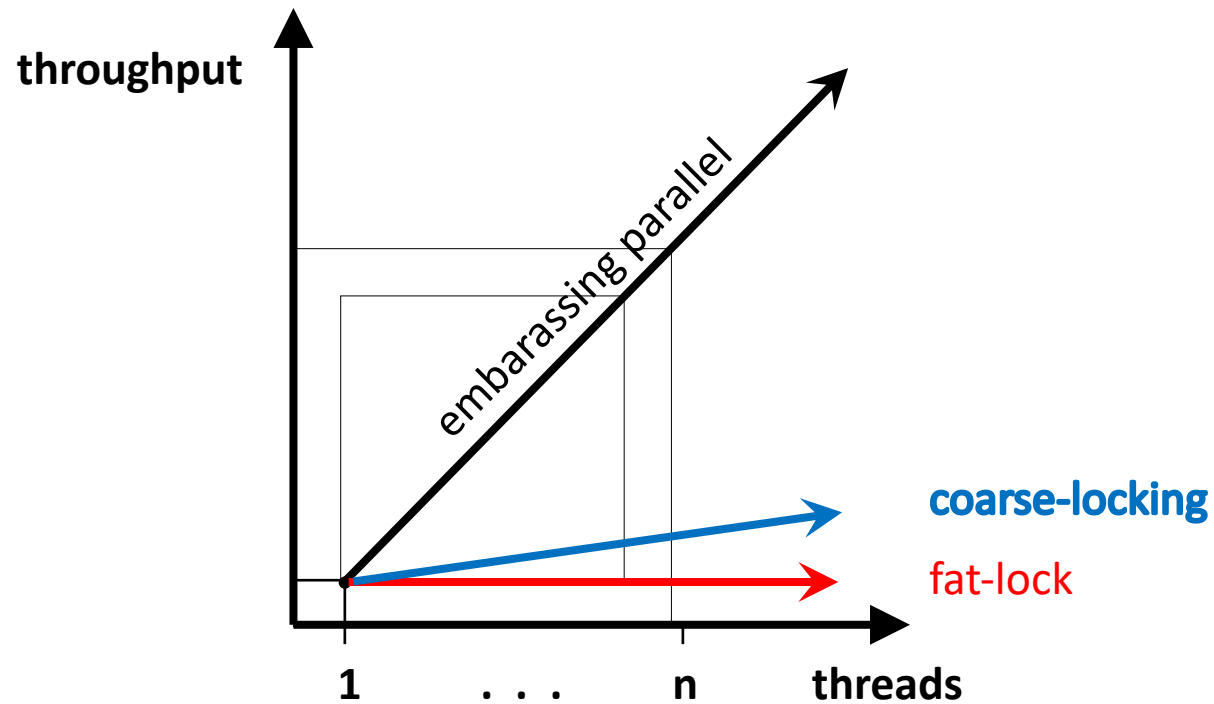
# Parallel Throughput



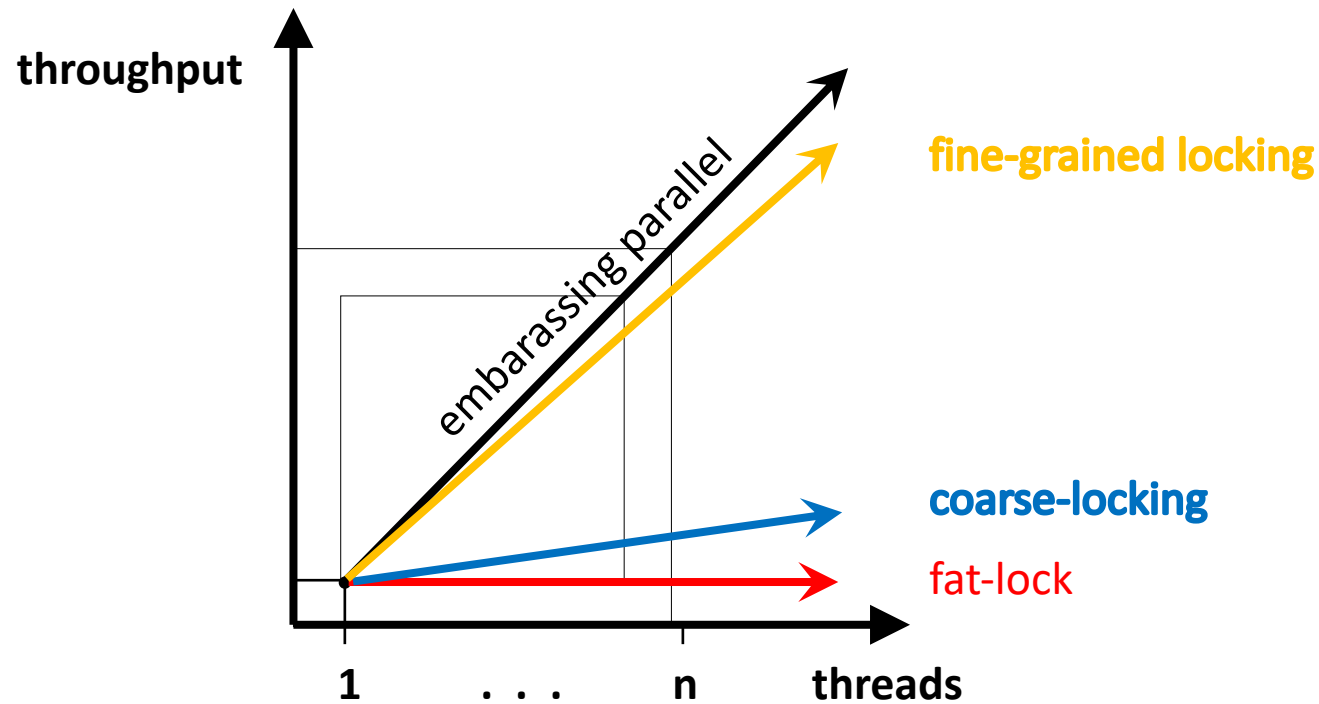
# Parallel Throughput



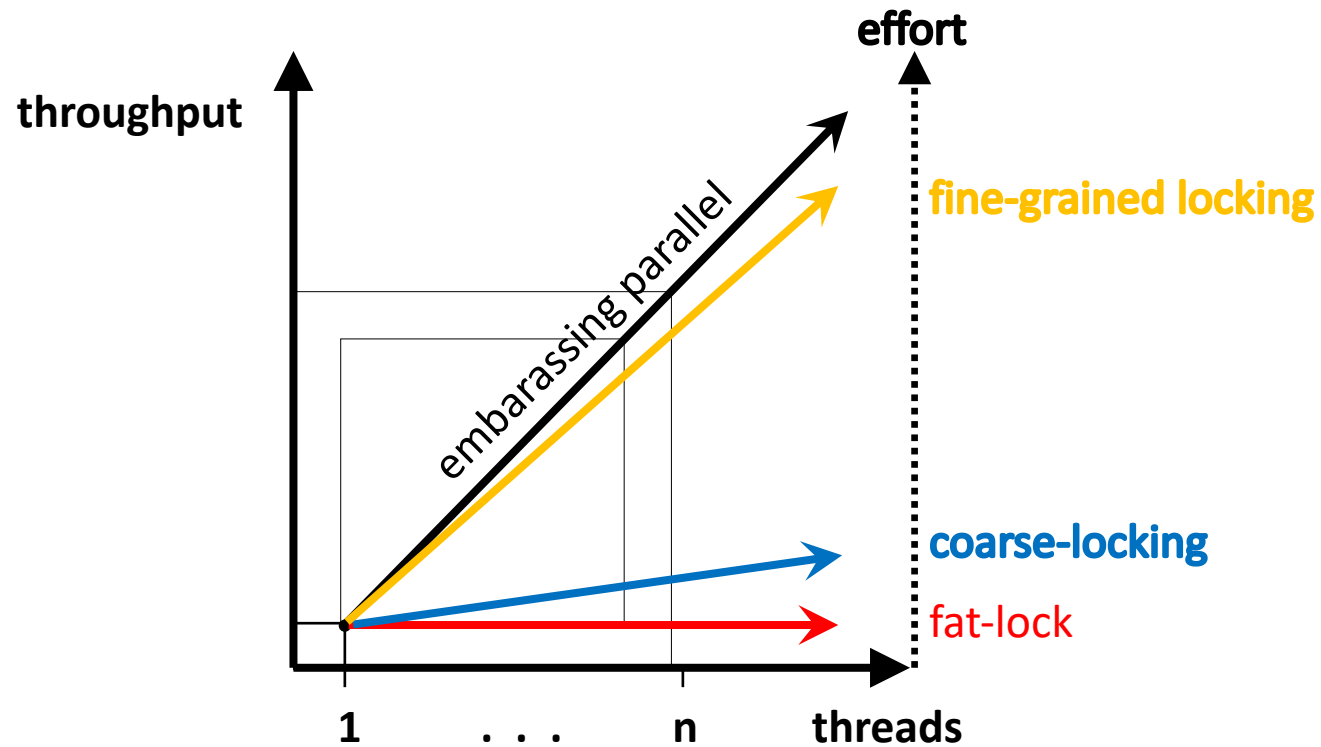
# Parallel Throughput



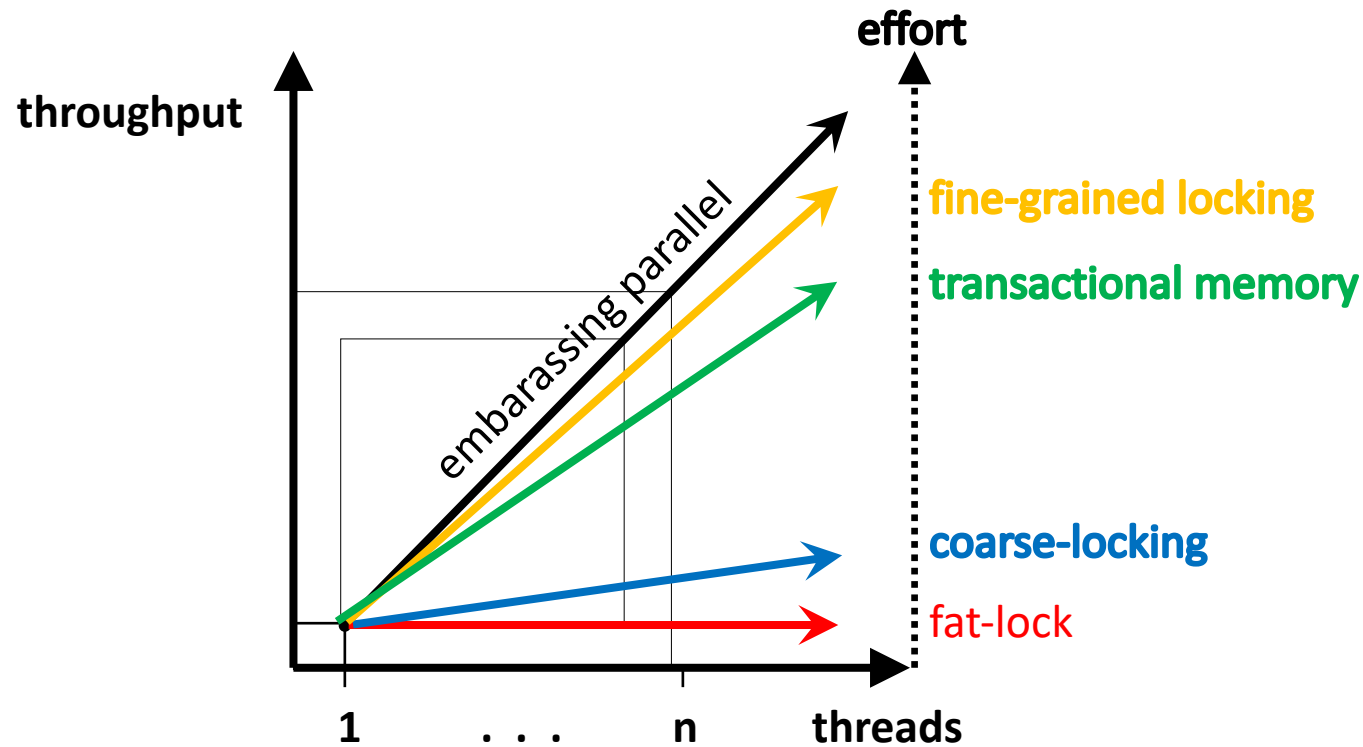
# Parallel Throughput



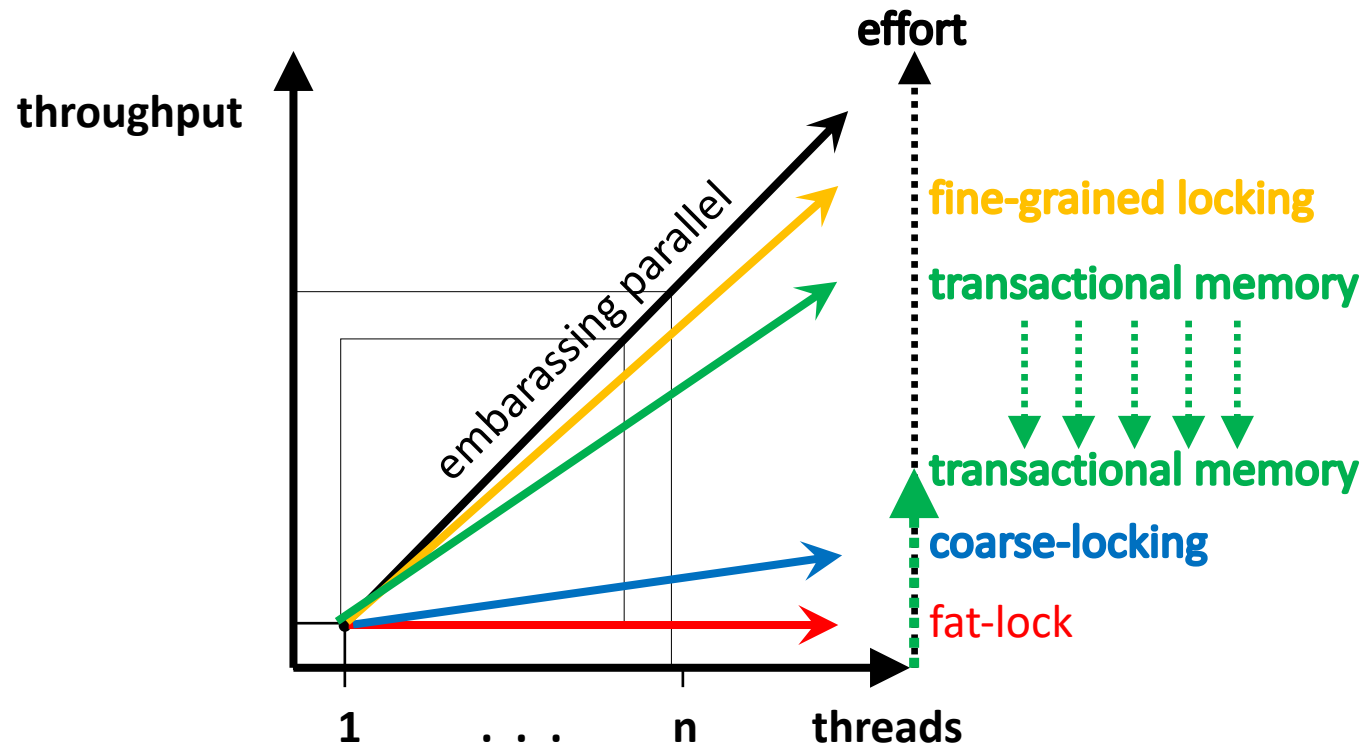
# Parallel Throughput



# Parallel Throughput



# Parallel Throughput





# What is

# Transactional Memory?



# What is transactional memory?

---

- Is an abstraction for simplifying concurrent programming that uses the concept of “transactions” to synchronize the accesses to shared data in main memory.

# Transactional computing

- If critical section locking is superfluous most of the time, aborts are rare
  - Typically threads manipulate different parts of the shared memory
  - Consider, e.g., web server serving pages for different users
- Optimistic approach
  - Instead of assuming that conflicts will happen in critical sections, assume they don't
  - Rely on conflict detection: abort and retry if necessary
- References
  - Lomet 1977, Herlihy & Moss 1993, Shavit & Touitou 1995, Herlihy et al. 2003

# Transactional memory

---

- “**atomic**”  $\approx$  “**transaction**”  $\Rightarrow$   
 $\Rightarrow$  *atomicity, consistency, isolation*

# Transactional memory

- “**atomic**”  $\approx$  “**transaction**”  $\Rightarrow$   
 $\Rightarrow$  *atomicity, consistency, isolation*
- Atomicity
  - One-or-nothing
- Consistency
  - Transactions take the program from one correct state into another correct state (assuming a bug-free program)
- Isolation
  - Transactions appear to execute alone in the system
    - i.e., with no interference from other concurrent transactions

# Transactional memory


- “**atomic**”  $\approx$  “**transaction**”  $\Rightarrow$   
 $\Rightarrow$  *atomicity, consistency, isolation*
- Atomicity
  - Commit: takes effect
  - Abort: effects rolled back
    - Usually retried
- Isolation
  - Serializable: any parallel execution of the transactions is equivalent one sequential execution of those same transactions — they always appear to happen in a one-at-a-time order

# The Brief History of (S)TM



# Locks:

## Insert in a Double Linked List



### Transactional memory

```
public void insertNode (node, precedingNode) {  
    atomic { //means "transaction"  
        node.prec = precedingNode;  
        node.succ = precedingNode.succ;  
        precedingNode.succ.prec = node;  
        precedingNode.succ = node;  
    }  
}
```

**Software or Hardware  
run-time manages the  
conflicting accesses!**



# Support for TM run-time(s)

- Hardware

- Sun Rock processor (abandoned)
- AMD HTM specification (in simulator)
- IBM BlueGene/Q (available)
- IBM Power8 (available)
- IBM zEC12 (available)
- Intel Haswell processor (available)
- Arm Processor V9 (available)

- Software

- Intel C++ compiler / GNU GCC 4.7+
- Haskell, Scala and Closure programming languages
- Java (annotations)

# TM Pros & Cons

- Advantages
  - Code blocks are simply marked as atomic/isolated
    - Less concerns about parallelism
  - Run-time ensure the ACI properties
    - Hide away synchronization issues from the programmer
- Disadvantages
  - Overhead
  - Successive collisions → repetitive restart → livelocks
  - Memory transactions can not contain certain operations (e.g., I/O operations)
  - Incompatibility with legacy code
  - ...

# Memory transaction life-cycle

---

1. Start
2. Access shared data (read / write)
3. (Attempt to) Commit
4. If commit fails, go to 1

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;

    curr = set → head;

    next = curr → next;
    while (next → val < val) do {

        curr = next;

        next = curr → next;
    }

    result = (next → val == val);

    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;

    curr = set → head;

    next = curr → next;
    while (next → val < val) do {

        curr = next;

        next = curr → next;
    }

    result = (next → val == val);

    return result;
}
```

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
lock(&set → head → lock);
    curr = set → head;
lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
unlock(&curr → lock);
        curr = next;
lock(&next → next → lock);
        next = curr → next;
    }
unlock(&curr → lock);
    result = (next → val == val);
unlock(&next → lock);
    return result;
}
```

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
transaction {
    curr = set → head;
lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
unlock(&curr → lock);
        curr = next;
lock(&next → next → lock);
        next = curr → next;
    }
unlock(&curr → lock);
    result = (next → val == val);
}
return result;
```

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
transaction {
    curr = set → head;

    next = curr → next;
    while (next → val < val) do {
        curr = next;
        next = curr → next;
    }

    result = (next → val == val);
}
    return result;
}
```



# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
    lock(&set → head → lock);
    curr = set → head;
    lock(&curr → next → lock);
    next = curr → next;
    while (next → val < val) do {
        unlock(&curr → lock);
        curr = next;
        lock(&next → next → lock);
        next = curr → next;
    }
    unlock(&curr → lock);
    result = (next → val == val);
    unlock(&next → lock);
    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
atomic {
    curr = set → head;

    next = curr → next;
    while (next → val < val) do {

        curr = next;

        next = curr → next;
    }

    result = (next → val == val);
}
    return result;
}
```

# Locks vs. Transactions

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;

    curr = set → head;

    next = curr → next;
    while (next → val < val) do {

        curr = next;

        next = curr → next;
    }

    result = (next → val == val);

    return result;
}
```

```
bool lk_contains(val) {
    int results;
    node_lk *prev, *next;
atomic {
    curr = set → head;

    next = curr → next;
    while (next → val < val) do {

        curr = next;

        next = curr → next;
    }

    result = (next → val == val);
}

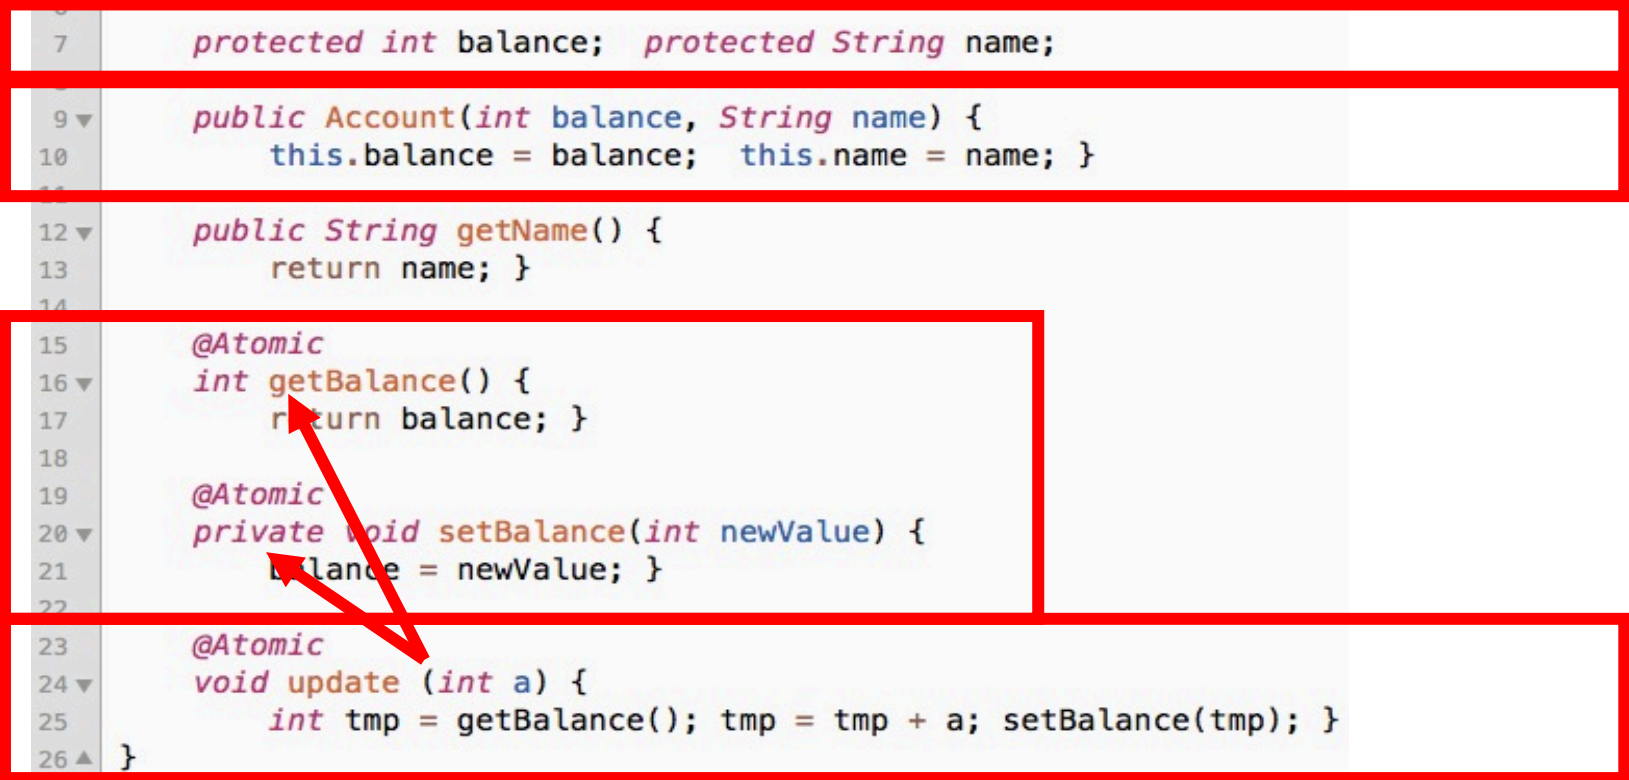
    return result;
}
```

# An example

```
1 package test.AccountTest;
2
3 public class Main {
4
5     static Account a;
6     public static void main(String[] args) {
7         //this will be accessed by both threads
8         a = new Account(0, "Account name");
9         new Update().start();
10        new Update().start();
11    }
12 }
```

# An example

```
1 package test.AccountTest;
2
3 import pt.moth.annotation.Atomic;
4
5 public class Account {
6
7     protected int balance; protected String name;
8
9     public Account(int balance, String name) {
10         this.balance = balance; this.name = name; }
11
12     public String getName() {
13         return name; }
14
15     @Atomic
16     int getBalance() {
17         return balance; }
18
19     @Atomic
20     private void setBalance(int newValue) {
21         balance = newValue; }
22
23     @Atomic
24     void update (int a) {
25         int tmp = getBalance(); tmp = tmp + a; setBalance(tmp); }
26 }
```



The diagram consists of three red rectangular boxes highlighting different parts of the code. The first box highlights the class declaration and the first two fields: `protected int balance; protected String name;`. The second box highlights the constructor and the `getName()` method. The third box highlights the `getBalance()`, `setBalance()`, and `update()` methods. Two red arrows originate from the `@Atomic` annotations on lines 15 and 19, pointing to the `return balance;` statement in the `getBalance()` method and the `balance = newValue;` statement in the `setBalance()` method, respectively.

# An example

```
1 package test.AccountTest;
2
3 import pt.moth.annotation.Atomic;
4
5 public class Account {
6
7     protected int balance; protected String name;
8
9     public Account(int balance, String name) {
10         this.balance = balance; this.name = name; }
11
12     public String getName() {
13         return name; }
14
15     @Atomic
16     int getBalance() {
17         return balance; }
18
19     @Atomic
20     private void setBalance(int newValue) {
21         balance = newValue; }
22
23     @Atomic
24     void update (int a) {
25         int tmp = getBalance(); tmp = tmp + a; setBalance(tmp); }
```

# An example

```
1 package test.AccountTest;
2
3 import pt.moth.annotation.Atomic;
4
5 public class Account {
6
7     protected int balance; protected String name;
8
9     public Account(int balance, String name) {
10         this.balance = balance; this.name = name; }
11
12     public String getName() {
13         return name; }
14
15     @Atomic
16     int getBalance() {
17         return balance; }
18
19     @Atomic
20     private void setBalance(int newValue) {
21         balance = newValue; }
22
23     @Atomic
24     void update (int a) {
25         int tmp = getBalance(); tmp = tmp + a; setBalance(tmp); }
```

```
1 package test.AccountTest;
2
3 public class Main {
4
5     static Account a;
6     public static void main(String[] args) {
7         //this will be accessed by both threads
8         a = new Account(0, "Account name");
9         new Update().start();
10        new Update().start();
11    }
12 }
```

```
1 package test.AccountTest;
2
3 import java.util.Random;
4
5 public class Update extends Thread {
6     public void run() {
7         while(true){
8             Random r = new Random();
9             int n = r.nextInt();
10            //Example.a.update(123);
11            Main.a.update(n);
12        }
13    }
14 }
```

# The END

---