

Can Software Transactional Memory make Parallel Programs Simple and Safe?

Ketil Malde

Institute of Marine Research

February 11, 2013

Outline

- Parallelism
- Concurrency Challenges
- Transactional Memory
- Genome Scaffolding
- Results
- Conclusions

Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

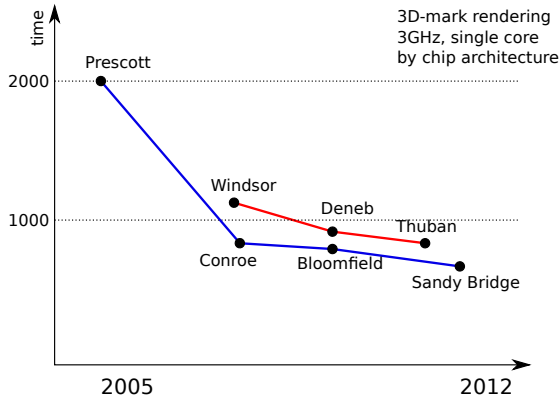
Transactional
Memory

Genome
Scaffolding

Results

Conclusions

CPU performance



source: tomshardware.com

Parallelism

Concurrency
Challenges

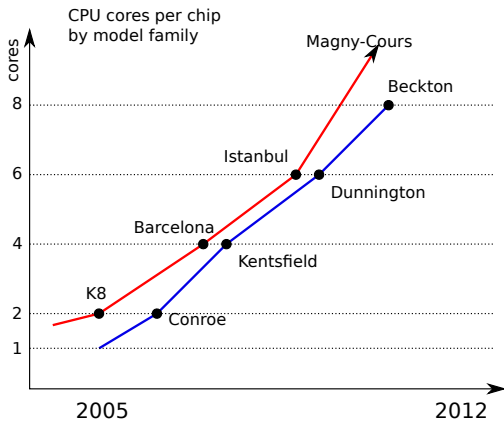
Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Increasing number of cores



source: wikipedia

Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions

2013: Octo-core phones?

Software
Transactional
Memory

Ketil Malde

Parallelism


Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions



Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Bank account example

```
deposit(acct, val) =  
  x <- readAcct(acct)  
  y = x+val  
  writeAcct(acct, y)
```


Bank account example

A: \$100

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x+10  
  writeAcct( A , y)
```

Bank account example

A: \$100

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x+10  
  writeAcct( A , y)
```

x=100

Bank account example

A: \$100

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x+10  
  writeAcct( A , y)
```

x=100 y=110

Banc account example

A: \$110

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x+10  
  writeAcct( A , y)
```

x=100 y=110

Race condition!

A: \$???

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x+10  
  writeAcct( A , y)
```

x=100 y=110

```
deposit( A , 20 ) =  
  x <- readAcct( A )  
  y = x+20  
  writeAcct( A , y)
```

x=100 y=120

Solution: Locking

```
deposit( A , 10 ) =  
  take_lock(A)  
  x <- readAcct( A )  
  y = x+ 10  
  writeAcct( A , y)  
  release_lock(A)
```

```
deposit( A , 20 ) =  
  take_lock(A)  
  x <- readAcct( A )  
  y = x+ 20  
  writeAcct( A , y)  
  release_lock(A)
```

Solution: Locking?

It is difficult to get right.

- ▶ Need to protect **all** shared data
- ▶ Deadlocks
- ▶ Livelocks
- ▶ *Locks prevent composability*

Composability

```
transfer(src,dst,amount) =  
  withdraw(src,amount)  
  deposit(dst,amount)
```


Composability

transfer(src,dst,amount) =
withdraw(src,amount)
deposit(dst,amount) ← Inconsistent state!

Composability

```
transfer(src,dst,amount) =  
  take_lock(src)  
  take_lock(dst)  
  withdraw(src,amount)  
  deposit(dst,amount)  
  release_lock(dst)  
  release_lock(src)
```

Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

**Transactional
Memory**

Genome
Scaffolding

Results

Conclusions

Software Transactional Memory

```
deposit( A , 10 ) =  
  x <- readAcct( A )  
  y = x + 10  
  writeAcct( A , y)
```

```
deposit( A , 20 ) =  
  x <- readAcct( A )  
  y = x + 20  
  writeAcct( A , y)
```

Software Transactional Memory

```
deposit( A , 10 ) =
```

```
BEGIN
```

```
x <- readAcct( A )
```

```
y = x+ 10
```

```
writeAcct( A , y)
```

```
COMMIT
```

```
deposit( A , 20 ) =
```

```
BEGIN
```

```
x <- readAcct( A )
```

```
y = x+ 20
```

```
writeAcct( A , y)
```

```
COMMIT
```

Software Transactional Memory

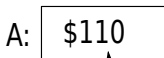
A: \$100

```
deposit( A , 10 ) =  
  BEGIN  
  x <- readAcct( A )  
  y = x + 10  
  writeAcct( A , y)  
  COMMIT  
  x=100  y=110
```

```
deposit( A , 20 ) =  
  BEGIN  
  x <- readAcct( A )  
  y = x + 20  
  writeAcct( A , y)  
  COMMIT  
  x=100  y=120
```

Software Transactional Memory

A: \$110



```
deposit( A , 10 ) =  
BEGIN  
x <- readAcct( A )  
y = x + 10  
writeAcct( A , y )  
COMMIT  
x=100 y=110
```

```
deposit( A , 20 ) =  
BEGIN  
x <- readAcct( A )  
y = x + 20  
writeAcct( A , y )  
COMMIT  
x=100 y=120
```

Software Transactional Memory

A: \$110

```
deposit( A , 10 ) =  
BEGIN  
x <- readAcct( A )  
y = x + 10  
writeAcct( A , y )  
COMMIT  
x=100 y=110
```

```
deposit( A , 20 ) =  
BEGIN  
x <- readAcct( A )  
y = x + 20  
writeAcct( A , y )  
COMMIT  
x=100 y=120
```

Retry

STM Performance Impact

- ▶ Logging data access
- ▶ Retrying transaction

Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

Transactional
Memory

**Genome
Scaffolding**

Results

Conclusions

Scaffolding

Given a set of *contigs*
And a set of *links*
Order and orient the contigs,
satisfying as many links as possible.

A graph problem:

- ▶ Contigs are nodes
- ▶ Link information make edges

Scaffolding

Given a set of *contigs*


And a set of *links*

Order and orient the contigs,
satisfying as many links as possible.


A graph problem:

- ▶ Contigs are nodes
- ▶ Link information make edges

Identify edge



Update graph



- ▶ Changes are local
- ▶ ..but unpredictable

→ Ideal for transactions?

Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions


Simple Implementation

Join contigs A and B iff:

B is the contig with most links from A
and

A is the contig with most links from B

Transactional Overhead



Parallelism

Concurrency
Challenges


Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Speedup



Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Caveats

Performance will be hurt by retried transactions.

Affected by:

- ▶ Number of threads/size of graph
- ▶ Size of subgraph affected by each transaction

Outline

Parallelism

Concurrency Challenges

Transactional Memory

Genome Scaffolding

Results

Conclusions

Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Conclusions

CPU time is precious — don't waste it on single-thread applications!

Programmer time is even more precious — don't waste it on primitive concurrency primitives!

Conclusions

CPU time is precious — don't waste it on single-thread applications!

Programmer time is even more precious — don't waste it on primitive concurrency primitives!

Conclusions

CPU time is precious — don't waste it on single-thread applications!

Programmer time is even more precious — don't waste it on primitive concurrency primitives!

The End

Software
Transactional
Memory

Ketil Malde

Parallelism

Concurrency
Challenges

Transactional
Memory

Genome
Scaffolding

Results

Conclusions

Thanks