

Today's agenda

- Act I: STMs — Why should we care?

The theoretical foundations

- Act II: STMs — How to build one?

The gory details



*“...we need to explore new techniques like **transactional memory** that will allow us to get the full benefit of all those transistors and map that into higher and higher performance.”*

Bill Gates, Businessman

Designing Transactional Memory Systems

Part I: Introduction

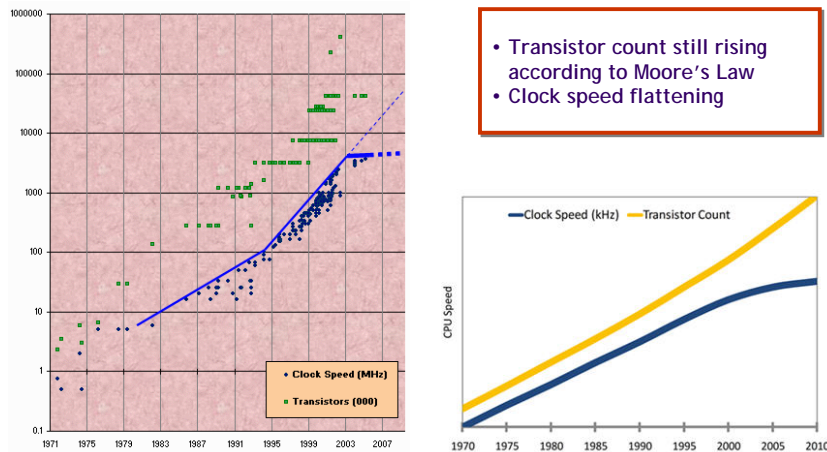
Pascal Felber
University of Neuchâtel
Pascal.Felber@unine.ch



Based on joint work with Christof Fetzer & Torvald Riegel
with slides borrowed from several other people

- **Part I:** Introduction
 - Why do we need STMs?
 - What do STMs provide?
 - A brief history of STMs
- **Part II:** Obstruction-free STMs
- **Part III:** Lock-based STMs

Moore's Law and CPU speed



3/17/2008

Transactional Memory: Part I — P. Felber

5

Multicores will be everywhere

- Multicores are the answer to keeping up with increasing CPU performance despite:
 - The **memory wall** (gap between CPU and memory speeds)
 - The **ILP wall** (not enough instruction-level parallelism to keep the CPU busy)
 - The **power wall** (higher clock speeds require more power and create thermal problems)
- Consequence:
 - Single-thread performance doesn't improve...
... but we can put more cores on a chip

3/17/2008

Transactional Memory: Part I — P. Felber

6

Multicores **are** everywhere

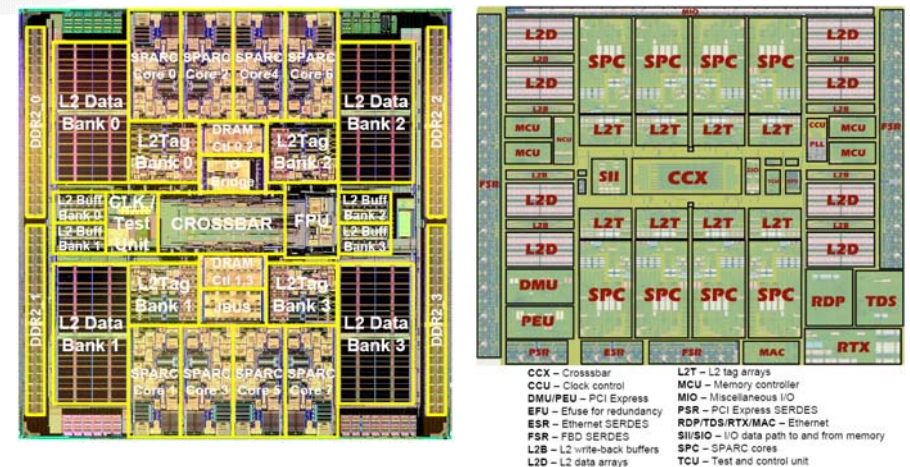
- **Dual-core** commonplace in laptops
- **Quad-core** in desktops
- **Dual quad-core** in servers
- All major chip manufacturers produce multicore CPUs
 - **SUN Niagara** (8 cores, 32 concurrent threads)
 - **Intel Xeon** (4 cores)
 - **AMD Opteron** (4 cores)
 - ...

3/17/2008

Transactional Memory: Part I — P. Felber

7

SUN's Niagara CPU2 (8 cores)

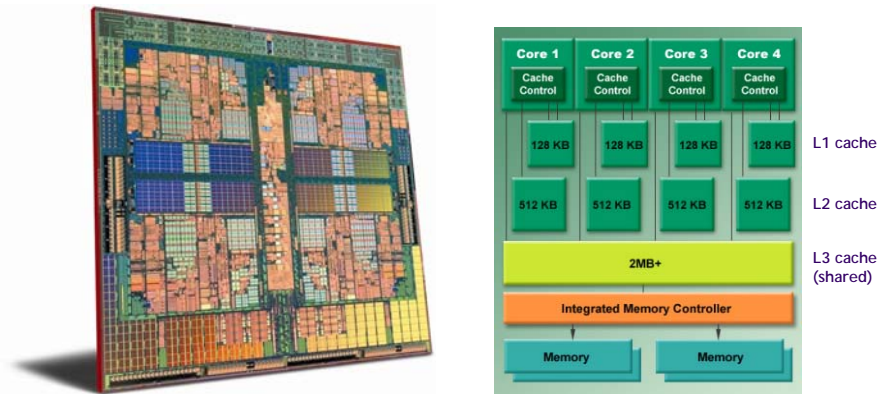


3/17/2008

Transactional Memory: Part I — P. Felber

8

AMD Opteron (4 cores)



3/17/2008

Transactional Memory: Part I — P. Felber

9

The “free ride” is over

- Cannot rely on CPUs getting faster in every generation
 - Utilizing more than one CPU core requires thread-level parallelism (TLP)
 - One of the biggest future software challenges: **exploiting concurrency**
 - Every programmer will have to deal with it
 - Affects HW/SW system architecture, programming languages, algorithms, ...
 - Concurrent programming is hard to get right
- ... better not hit the **productivity wall**

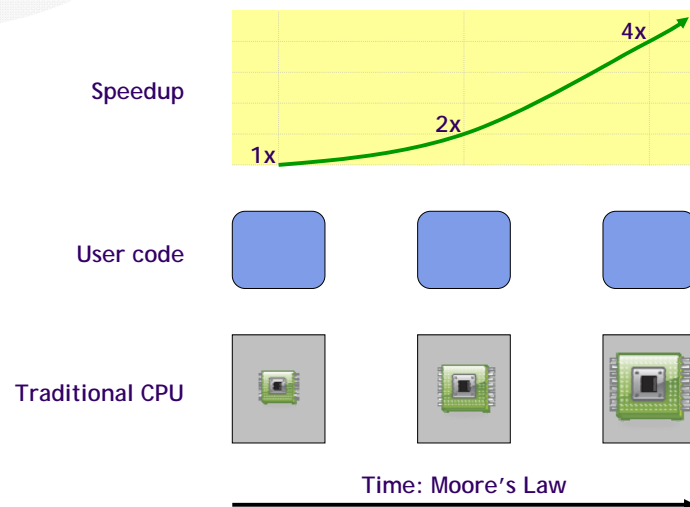
3/17/2008

Transactional Memory: Part I — P. Felber

10

Based on slide by Herlihy & Shavit

Traditional scaling process




3/17/2008

Transactional Memory: Part I — P. Felber

11

Based on slide by Herlihy & Shavit

Multicore scaling process

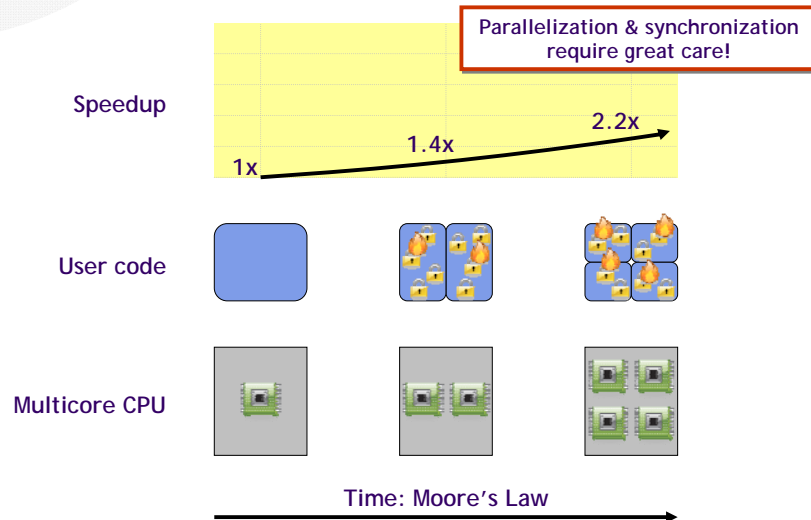


3/17/2008

Transactional Memory: Part I — P. Felber

12

Real-world scaling process



Concurrent programming is hard

- Hard to make **correct** and **efficient**
 - We need to exploit parallelism
- The human mind tends to be sequential
 - Concurrent specifications
 - Non-deterministic executions
- What about races? deadlocks? livelocks? starvation? fairness?
 - Need synchronization (correctness)...
 - ... but not too much (performance)

Parallelization (1)

- **Data parallelism**
 - Split data into partitions
 - Let threads process partitions
 - Threads join after finishing their work
 - **Example:** image processing, sequencing, etc.
- Good data partitioning is difficult (correctness, load balancing, ...)
- Synchronization required
 - Join
 - Load balancing during runtime

Parallelization (2)

- **Pipeline parallelism**
 - Split work into phases
 - Let each thread work on jobs in one of the phases
 - **Example:** event stream processing
input → parse → process → format → output
- Not every program has such phases, some phases are longer than others
- Synchronization required
 - One phase's results become next phase's input
 - Complex access patterns

Parallelization (3)

- **Speculative/optimistic parallelism**
 - Just execute concurrent jobs
 - Assumption: most jobs do not conflict
 - Partitioning or phases not required but possible
 - **Example:** event handlers, application servers, graph algorithms, etc.
- There must be little contention on shared data structures
- Synchronization required
 - Every access can potentially interfere with accesses from another thread

3/17/2008

Transactional Memory: Part I — P. Felber

17

Shared memory synchronization

- Cores share main memory
 - Some hardware instructions are (or can be made) atomic: `inc`, `dec`, `cmpxchg`, ...
- Loads/store usually atomic, not necessarily ordered (no sequential consistency!)
 - Must use memory barriers to enforce order
- Current state of concurrent programming:
 - Use locks built from these instructions
 - Build concurrent (non-blocking) algorithms from these instructions

3/17/2008

Transactional Memory: Part I — P. Felber

18

Based on slide by Grossman

Why transactions?

```
void deposit(...) { synchronized(this) { ... } }
void withdraw(...) { synchronized(this) { ... } }
int balance(...) { synchronized(this) { ... } }

void transfer(account from, int amount) {

    if (from.balance() >= amount) {
        from.withdraw(amount);
        this.deposit(amount);
    }

}
```

No concurrency control: race!

3/17/2008

Transactional Memory: Part I — P. Felber

19

Based on slide by Grossman

Why transactions?

```
void deposit(...) { synchronized(this) { ... } }
void withdraw(...) { synchronized(this) { ... } }
int balance(...) { synchronized(this) { ... } }

void transfer(account from, int amount) {

    synchronized(this) {
        if (from.balance() >= amount) {
            from.withdraw(amount);
            this.deposit(amount);
        }
    }

}
```

Race!

3/17/2008

Transactional Memory: Part I — P. Felber

20

Why transactions?

```
void deposit(...) { synchronized(this) { ... } }
void withdraw(...) { synchronized(this) { ... } }
int balance(...) { synchronized(this) { ... } }
```

```
void transfer(account from, int amount) {
    synchronized(this) {
        synchronized(from) {
            if (from.balance() >= amount) {
                from.withdraw(amount);
                this.deposit(amount);
            }
        }
    }
}
```

Deadlock!

Atomic blocks

```
void deposit(int x) {
    synchronized(this) {
        int tmp = balance;
        tmp += x;
        balance = tmp;
    }
}
```

Lock acquire/release

```
void deposit(int x) {
    atomic {
        int tmp = balance;
        tmp += x;
        balance = tmp;
    }
}
```

(As if) no interleaved computation

Easier-to-use primitive
(but harder to implement)

Atomic blocks

```
void deposit(...) { atomic { ... } }
void withdraw(...) { atomic { ... } }
int balance(...) { atomic { ... } }
```

```
void transfer(account from, int amount) {
    if (from.balance() >= amount) {
        from.withdraw(amount);
        this.deposit(amount);
    }
}
```

No concurrency control: race!

Atomic blocks

```
void deposit(...) { atomic { ... } }
void withdraw(...) { atomic { ... } }
int balance(...) { atomic { ... } }
```

```
void transfer(account from, int amount) {
    atomic {
        if (from.balance() >= amount) {
            from.withdraw(amount);
            this.deposit(amount);
        }
    }
}
```

Correct and enables parallelism!

Why STM?

- Transactions: a simple paradigm
 - A sequence of instructions, executed atomically
- Software transactions are good for:
 - Software engineering (simple programming, avoid races & deadlocks, composability)
 - Performance (when no conflict, high parallelism and no locking overhead)

A “universal” synchronization construct

Don't care how transactions are implemented!

Composability

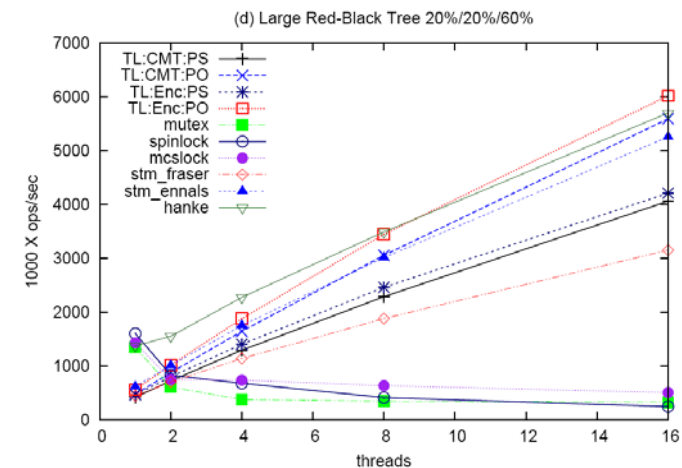
- Ability to build large, complex systems from small, simple pieces
 - Enables divide-and-conquer strategy
- Locks are **not** composable
 - Must be exposed and use “compatible” strategies
- Custom concurrent algorithms are typically **not** composable
- Composability challenges
 - Must preserve correctness
 - Should not hamper performance
- STM **is** composable

Can we make STM fast?

- Typically, problem-specific handcrafted algorithms are **more efficient**...
 - Whether lock-free or using fine-grained locking
 - Programmers can use knowledge of data flow relationships to control contention
- ... but STM is **simpler and safer** to use...
 - Programmer does not need to reason about concurrency
- ... and there is room for **optimism**
 - Performance can be as good or better for complex data structures

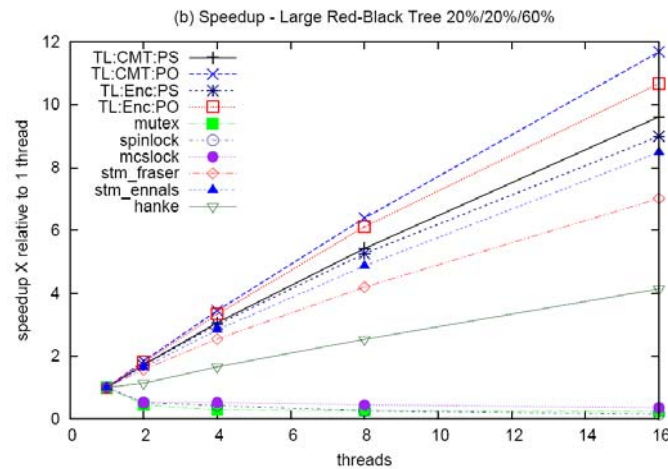
STM performance [Dice & Shavit]

STM can be as efficient as handcrafted lock-based implementation!



STM scaling [Dice & Shavit]

STM scales better than handcrafted lock-based implementation!



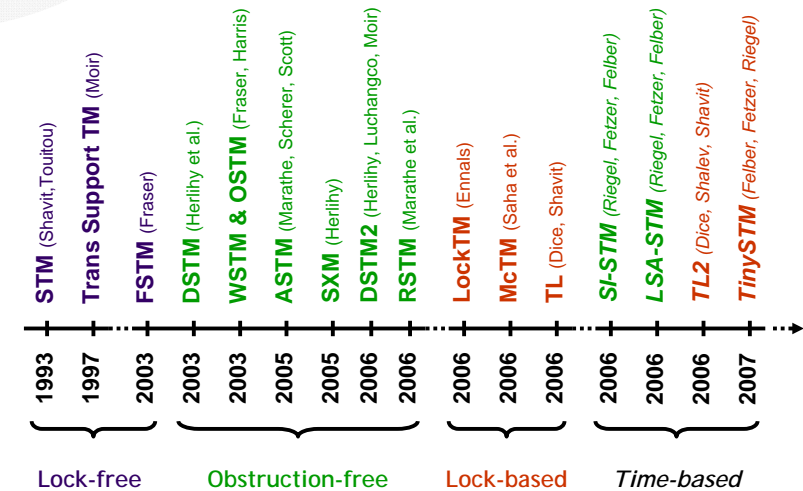
3/17/2008

Transactional Memory: Part I — P. Felber

29

Based on slide by Shalev & Shavit

A brief (partial) history of STM



3/17/2008

Transactional Memory: Part I — P. Felber

30

Conclusion (Part I)

- **Multicores are here!**
 - No more scalability for free
 - We need the right tools to exploit their power
- **Concurrent programming is hard**
 - Synchronization necessary for correctness, parallelism for efficiency
- STM exploits **disjoint access parallelism**
 - Execute sequences of operations atomically (don't care how this is done!)
 - Optimistic concurrency control
 - Transactions are composable

3/17/2008

Transactional Memory: Part I — P. Felber

31