

# Transaction Concurrency Control via Dynamic Scheduling Based on Static Analysis

Paweł T. Wojciechowski, Konrad Siek  
Poznań University of Technology



Bern, 10 April 2012

Transactional Memory (TM) traditionally employs optimistic concurrency control + contention management.

Advantages of fully pessimistic approaches were recently recognized (Shavit, Matveev WTTM '11): I/O and system calls, debugging, ...

A lot fewer results are about *distributed TMs*.

- ▶ **optimistic** locally scoped transactions + replicated objects => Paxos STM (Euro-TM WDTM '12),
- ▶ **pessimistic** distributed transactions on remote objects => Atomic RMI (this talk).

## EXAMPLE: “WITHDRAW-DEPOSIT” TRANSACTION


Below is a program in Atomic RMI, where a and b are remote objects (in the sense of Java RMI):

```
Transaction t = new Transaction( registry );


t.start();
int balance = a.getBalance();
if ( balance < sum ) {
    t.rollback();
} else {
    a.withdraw( sum );
    b.deposit( sum );
    t.commit();
}
```

# PESSIMISTIC CONCURRENCY BY VERSIONING<sup>1</sup>

Basic Versioning Algorithm:



Supremum Versioning Algorithm:




Static (a priori) scheduling:  $T_2$  is temporarily blocked by  $T_1$ .  
 $T_3$  is **not** blocked since it accesses other objects.

Both schedules are serializable but  $C'_{max} < C_{max}$ .

<sup>1</sup>Wojciechowski *et al.* IPDPS '04, PDPD '05

# TRANSACTION COMMIT AND ROLLBACK (SVA)



Transaction thread is blocked till transaction commit ( $\diamond$ ).

$T_1$ 's rollback causes  $T_2$  and  $T_4$  to rollback  $\Rightarrow$

$T_2$ 's commit must be **deferred**.

$T_2$ 's rollback must cause  $T_4$  to rollback.

$T_3$  is not affected since it accesses other objects.

When to release an object held by a transaction?

We use 3 methods:

1. The programmer releases an object (simple but not safe!).
2. **Supremum analysis**: Release an object if # of object calls reached an upper bound (supremum  $S$ ).
3. **Last-use analysis**: If  $S$  cannot be inferred, release an object if no more calls can be made by a client (work-in-progress).

We developed a precompiler that instruments Java code after analyzing its translation to Jimple (more suitable than Java bytecode).

## EXAMPLE CODE AFTER PRECOMPILE ROUND

Below is our example program, modified by the precompiler:

```
Transaction t = new Transaction( registry );
```

```
a = t.accesses(a, 2); // upper bound = 2
```

```
b = t.accesses(b, 1); // upper bound = 1
```

```
t.start();  
int balance = a.getBalance();  
if ( balance < sum ) {  
    t.rollback();  
} else {  
    a.withdraw(sum);  
    b.deposit(sum);  
    t.commit();  
}
```



# SUPREMUM VS. LAST-USE ANALYSIS

The object last-use analysis applied:

```
transaction.start();  
for (i=0; i < n; i++) {  
    a.m();  
    b.m(); // a unnecessary blocked in the last loop  
}  
::release a, b;  
transaction.commit();
```

The object call supremum analysis applied:

```
transaction.start();  
for (i=0; i < n; i++) {  
    a.m(); // a will be released in nth iteration  
    b.m(); // b will be released in nth iteration  
}  
transaction.commit();
```

# OBJECT RELEASE BASED ON LAST-USE ANALYSIS

Clients on nodes  $i$  ( $i = 2..n$ ):

```
transaction.start();  
a.m();  
:: release a;  
transaction.commit();
```

A method of a remote object  $a$  on node 1:

```
m() {  
    :: while (object blocked) {};  
    ...  
    return res;  
}
```

Release of object  $a$  by every client requires a network message.

Clients on nodes  $i$  ( $i = 2..n$ ):

```
transaction.start();  
a.m();  
transaction.commit();
```

A method of a remote object  $a$  on node 1:

```
m() {  
  :: while (ticket(i) - obj_count > supremum(i)) {};  
  ...  
  :: obj_count ← obj_count + 1  
  return res;  
}
```

Release of object  $a$  is a local operation (no network overhead).

1. On transaction start-up, request a globally unique ticket (required by static scheduling of object calls).
2. Sort the queue of ticket requests, either:
  - ▶ long transactions first, or
  - ▶ short transactions first.
3. Use a simple heuristic: the transaction length is proportional to statically inferred # of object calls.

Designed to check if experimental results match a theoretical schedule (a larger benchmark is an ongoing work):


- ▶ four remote objects a, b, c, and d,
- ▶ each object has one method “sleep 1 second and return”,
- ▶ 4-24 concurrent transactions spawned on 8 network nodes at the same time:
  - ▶ 50% **short transactions** (3 sec.): calls of a, b, and c,
  - ▶ 50% **long transactions** (6 sec.): calls of c, d, c, d, a, and b.

Evaluation environment:

- ▶ cluster of nodes: quad-core Xeon X3230 CPU with 4 GB RAM,
- ▶ 1Gb network.


# EXPERIMENTAL RESULTS

## Mean Flow Time




8 Transactions

## Throughput




8 Transactions

## Mean Flow Time



24 Transactions

## Throughput



24 Transactions

# CONCLUSIONS AND FUTURE WORK

- ▶ Pessimistic concurrency by versioning outperforms global lock.
- ▶ Static analysis can make this approach efficient and safe.
- ▶ Dynamic scheduling of transactions can further improve it for almost null cost at runtime (done in background).
- ▶ Upper bounds on calls are used to estimate transaction length.

Future work:

- ▶ Problems of distributed deadlock and partial faults.

Project page:

- ▶ <http://www.it-soa.eu/atomicrmi>