

1 Transactional Memory

(Notes from this section are based on the slides for the book “The Art of Multiprocessor Programming” by Herlihy and Shavit.)

Multi-core technology has brought concurrent programming back to the agenda of developers and researchers.

- Clock speeds are stagnant, which implies that time no longer cures software bloat.
- Applications must exploit concurrency from multi-core architectures to become faster.

In the beginning of this course, we proposed the use of techniques based on locking to address the issue of synchronizing concurrent access to shared data. Such techniques suffer from some problems:

- Overhead can be high, especially if applications resort to coarse-grained locking to avoid synchronization errors.
- They are not robust to thread/process failure (what happens when a thread fails while holding a lock?)
- They rely on conventions to avoid deadlocking. As we saw, one way to avoid deadlocks was to establish an ordering in which locks can be acquired. This leads to complex conventions about lock acquisition (many such examples can be found in the Linux kernel).
- They can be difficult to use (as you probably noticed from your assignments).
- Importantly, they are not composable. E.g., if I am given a thread-safe library that has operations to add and remove an element from a finite queue, how do I implement a thread-safe operation that moves an element from one queue to another, but does not change anything if the target buffer is full? Here’s an incorrect attempt:

```
xferVal(q1,q2) {
    q1.qlock.lock();
    q2.qlock.lock();
    if (source is non-empty and dest non-full) {
```

```

        x=deq(q1);
        enq(q2,x);
    }
    q1.qlock.unlock();
    q2.qlock.unlock();
}

```

This not only breaks the abstraction of how these things are synchronized, but also may incur in deadlock if one thread calls this with queues A,B and another with B,A.

Currently we are seeing many proposals for replacing locking with a transactional API (or a so-called *Transactional Memory*).

These transactions are similar to the ones we have seen before, but do not require the “permanence” property: data modified by a transaction that committed does not have to persist after a system failure.

Programming model is similar to the one we have seen before. A block of code is declared to be part of a transaction, which implies it is:

- Atomic – it can *commit* and its changes take effect, or *abort* in which case its changes are rolled back and (in most proposals) the transaction is retried automatically.
- Linearizable (a.k.a. Isolation property) – Transaction appear to have executed in sequence (one at a time), and that sequence is consistent with the real-time order of the execution of the program.

Simple example:

```

atomic {
    if (source is empty or dest full)
        abort(); // undoes possible previous effects
    x=deq(q1);
    enq(q2,x);
}

```

(More complex primitives are available, such as ability for programmer asking for retry a txn but only when one of the read variables is changed, etc.)

Also, integration with existing language mechanisms is still open.

Implementations

Example of an important implementation, DSTM by Herlihy et al., one of the initial implementations that influenced many of the current proposals.

DSTM stands for Dynamic STM. Transactions modify private copies of data (without requiring any synchronization), and replace original versions upon commit. At commit time, checks are made to see if conflicts occurred, in which case one of the conflicting transactions must abort (and automatically restarted). When data items are first read within a txn, they're added to the txn's read set (a set of pointers to objects read by the txn which is maintained by the system object that describes each transaction).

When data items are first modified within a txn, system creates a copy of item, only visible to this txn until committed. For this purpose, the system maintains current, read-only version and a modifiable version used by a pending transaction. Validation determines if a transaction has to abort, in which case the system will automatically restart it. Conflicts are detected either (1) when transactions first access a data item that is being modified by another transaction (early detection) or (2) upon commit when a transaction validates the set of data items it accessed to check if any of these were concurrently written by other transactions (late detection).

Validation of read/write conflicts in DSTM works by scanning read set and ensuring that for all the objects that were read, no concurrently executing transaction modified that object (this is done by checking if there are other versions maintained by pending transactions or if the version that was read has been superseded by a new version). DSTM does this validation both at commit time and also when it first reads an object, to avoid a transaction to continue executing when it read an inconsistent state. The other step for validation to work is to avoid write-write conflicts, where two txns modify the same object. This is achieved by performing this check whenever a transaction first modifies an object and aborting one of the conflicting transactions if needed.

If a transaction commits DSTM replaces all modified objects atomically with the new versions (using several data structures and special instructions we discussed before such as CAS or LL/SC).

Open issues

There are many open issues with respect to other aspects of the programming model, e.g.:

- How to deal with I/O? If a transaction sends a message and then aborts, how to deal with this? If I/O is also transactional (e.g. file system) it's easy to deal by aborting that transaction as well, but for other types solution might be to buffer all I/O until transaction commits. Still, this does not address the problem of a transaction that outputs a string to a terminal and subsequently waits for user input (within the same transaction).
- How to integrate with languages that support exceptions? What if an exception is thrown inside a transaction?
- What are the semantics for nested transactions (transactions that are started inside a transaction)?
- How to integrate transactional and non-transactional code? Is the non-transactional code allowed to see a non-committed change?
- How to handle conditional synchronization? Suppose a producer-consumer where consumer is waiting for item to be produced. One possibility for consumer to read buffer and abort if empty, but this amounts to busy waiting. Thus there have been proposals to extend txns with a guard that prevents a txn from starting until a predicate is true. Other proposals exist.

STM have considerable overhead, which can be a show-stopper. Fortunately hardware can help, and some of the many proposals for providing hardware support for transactional memory (called Hardware Transactional Memory or HTM) are starting to be implemented. Exact details vary considerably according to proposals, but they tend to work only for small transactions, and fall back to software when a transaction accesses lots of data.