

Durability and Contention in Software Transactional Memory

Masterstudium:

Software Engineering & Internet Computing

Michael Schröder

Technische Universität Wien
Institut für Computersprachen
Arbeitsbereich: Programmiersprachen und Übersetzer
Betreuer: Univ.-Prof. Dr. Jens Knoop
Mitbetreuerin: Assoc. Prof. Gabriele Keller, UNSW

Software Transactional Memory

Software Transactional Memory (STM) vastly simplifies concurrent programming by grouping memory operations into atomic blocks. The following Haskell function increments a transactional variable and returns its previous contents:

```
inc v = do x ← readTVar v
        writeTVar v (x + 1)
        return x
```

To perform an STM computation and make its effects visible to the system, the function `atomically :: STM a → IO a` is used:

```
atomically (inc v)
```

Problem 1: Durability



Problem 2: Contention



Manipulating memory using STM is easy, but persisting those manipulations in a transactionally safe way is impossible.

Two obvious but unsafe ways of trying to add durability to STM:

1. `do x ← atomically m`
`serialize x`

Problem: serialization might fail after transaction committed


2. `atomically $ do x ← m`
`unsafeIOToSTM (serialize x)`

Problem: transaction might abort or retry after serialization

Many standard data structures, when used in a transactional setting, cause unreasonably high numbers of conflicts.

Consider `TVar (HashMap a)`:

- ▶ any change to the container invalidates all other transactions
- ▶ but we should only care about the subset relevant to our transaction
- ▶ if transaction **A** updates element k_1 and transaction **B** deletes element k_2 and if $k_1 \neq k_2$, then there should be no conflict



STM Finalizers

Introduce a new STM primitive

```
atomicallyWithIO :: STM a → (a → IO b) → IO b
```

which is like `atomically`, but additionally takes a *finalizer* — an I/O action that can depend on the result of the STM computation.

The finalizer is combined with the STM transaction such that:

1. The finalizer is only run if the transaction is guaranteed to commit.
2. The transaction only commits if the finalizer finishes successfully.

$$\frac{M; \Theta, \{\} \Rightarrow \text{return } N; \Theta', \Delta' \quad F N; (\Theta \cup \Delta') \Rightarrow \text{return } P; \hat{\Theta} \quad (\text{ARET})}{\mathbb{P}[\text{atomicallyWithIO } M F]; \Theta \rightarrow \mathbb{P}[\text{return } P]; \Theta' \cup \hat{\Theta}}$$

$$\frac{M; \Theta, \{\} \Rightarrow \text{return } N; \Theta', \Delta' \quad F N; (\Theta \cup \Delta') \Rightarrow \text{throw } P; \hat{\Theta} \quad (\text{ATHROW2})}{\mathbb{P}[\text{atomicallyWithIO } M F]; \Theta \rightarrow \mathbb{P}[\text{return } P]; \hat{\Theta}}$$


the design is formalized by an operational semantics

Transactional Trie

The **transactional trie** is a new contention-free data structure, specifically tailored to the needs of transactional concurrency.

- ▶ based on the lock-free concurrent trie
- ▶ uses `unsafeIOToSTM` to perform atomic compare-and-swap operations independently of the surrounding STM transaction
- ▶ to preserve safety, leaves are stored as `TVar (Maybe a)`
- ▶ once a leaf is added, it is never removed

empirical evaluation of 200 000 random transactions on an Amazon EC2 C3 extra-large instance with 16 cores



- ▶ eliminates all spurious conflicts
- ▶ up to 8 times faster and using almost 10 times less memory
- ▶ <http://hackage.haskell.org/package/ttrie>

- ▶ durability is now trivially possible:

```
atomicallyWithIO m serialize
```

- ▶ more generally, finalizers enable interactive transactions
- ▶ potential foundation for a distributed STM
- ▶ <http://github.com/mcschroeder/ghc>



An example application using both finalizers and transactional tries to build a lightweight database framework on top of STM is available at <http://github.com/mcschroeder/social-example>