

Correctness of an STM Haskell Implementation

Manfred Schmidt-Schauß, David Sabel

Goethe-University, Frankfurt, Germany

ICFP '13, Boston, USA

Software Transactional Memory (STM)

- treats **shared memory** operations as **transactions**
- provides **lock-free** and **very convenient** concurrent programming
- requires an **implementation** that **correctly executes** the transactions

STM Haskell

- STM library for Haskell
- introduced by [Harris et.al](#), PPOPP'05
- uses Haskell's **strong type system** to distinguish between
 - software transactions,
 - functional code, and
 - IO-computations

Transactional Variables:

`TVar a`

Primitives to form STM-transactions `STM a`:

```
newTVar      :: a -> STM (TVar a)
readTVar     :: TVar a -> STM a
writeTVar    :: TVar a -> a -> STM ()

return      :: a -> STM a
(>>=)      :: STM a -> (a -> STM b) -> STM b

retry       :: STM ()
orElse      :: STM a -> STM a -> STM a
```

Executing an STM-transaction:

```
atomically :: STM a -> IO a
```

Semantics: the transaction-execution is

- **atomic:** all or nothing, effects are indivisible, and
- **isolated:** concurrent evaluation is not observable

Issues:

- Is an STM implementation **correct**?
- What does **correctness** mean?


Several correctness notions have been suggested

e.g. Guerraoui & Kapalka, PPOPP'08

- linearizability, serializability, recoverability, opacity, ...
- Most of these notions are **properties on the trace** of read-/write accesses on the transactional variables.

Our approach is **different**: “**semantic** approach”

Two program calculi for STM Haskell:



Correctness: The implementation fulfills the specification

➡ ψ is semantics reflecting

Adapted from the CHF-calculus (S.& Schmidt-Schauß: PPDP'11, LICS'12)

Processes:

$P_i \in Proc ::= P_1 \mid P_2 \mid \nu x.P \mid \underbrace{\langle u \mid x \rangle \Leftarrow e}_{\text{Concurrent future } x \text{ with identifier } u \text{ evaluates } e} \mid \overbrace{x \mathbf{t} e}^{\text{TVar } x \text{ with content } e}$

Expressions:

$e_i \in Exp ::= x \mid \lambda x.e \mid (e_1 e_2) \mid (c e_1 \dots e_{\text{ar}(c)})$
 $\mid \text{seq } e_1 e_2 \mid \text{letrec } x_1 = e_1, \dots, x_n = e_n \text{ in } e$
 $\mid \text{case}_T e \text{ of } \text{alt}_{T,1} \dots \text{alt}_{T,|T|}$
 where $\text{alt}_{T,i} = (c_{T,i} x_1 \dots x_{\text{ar}(c_{T,i})} \rightarrow e_i)$

$\mid \text{return}_{\text{IO}} e \mid e_1 \gg_{\text{IO}} e_2 \mid \text{future } e$
 $\mid \text{atomically } e \mid \text{return}_{\text{STM}} e \mid e_1 \gg_{\text{STM}} e_2$
 $\mid \text{retry} \mid \text{orElse } e_1 e_2$
 $\mid \text{newTVar } e \mid \text{readTVar } e \mid \text{writeTVar } e$

extended λ -calculus
 IO and STM

Monomorphic type system

Operational Semantics:

- **Call-by-need** “small-step” reduction \xrightarrow{SHF} , several rules, e.g.

(fork) $\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{future } e] \xrightarrow{SHF} \nu z, u'. (\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{return}_{\text{IO}} z] \mid \langle u'\lambda z \rangle \Leftarrow e)$

- **Big-step rule** for transactional evaluation:

$$\frac{\mathbb{D}_1[\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{atomically } e]] \xrightarrow{SHFA,*} \mathbb{D}'_1[\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{atomically } (\text{return}_{\text{STM}} e')]]}{\mathbb{D}[\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{atomically } e]] \xrightarrow{SHF} \mathbb{D}'[\langle u\lambda y \rangle \Leftarrow \mathbb{M}[\text{return}_{\text{IO}} e']]}$$

where \xrightarrow{SHFA} are small-step rules for transactional evaluation

- Enforces **sequential evaluation** of transactions
 - ➡ atomicity and isolation obviously hold
- Rule application is **undecidable!**

Extensions w.r.t. *SHF*:

- **local** and **global** TVars:
 - u **tls** S = Stack of **thread-local** TVars
 - x **tg** e u g = **global** TVar, where
 - u is a locking label (unlocked / locked by thread u)
 - g is a list of thread identifiers (the **notify list**)
- threads may have a **transaction log**: $\langle u\lambda y \rangle \xleftarrow{T,L;K} e$
 T, L, K are (stacked) lists storing information about created, read, and written TVars
- ...

Stacks are necessary for rollback during **nested orElse**-evaluation

Operational semantics:

- **true small-step** reduction \xrightarrow{CSHF}
- **concurrent** evaluation of STM transactions
- all rule applications are **decidable**

Transaction execution (informally):

- all read/writes are **logged** and performed on **local** TVars
- during the first readTVar-operation of thread u on TVar x :
 u is **added** to the **notify list** of TVar x
- commit phase
 - 1 **lock** global TVars
 - 2 **send a retry** to all threads in the **notify lists** of to-be-written TVars (= **conflicting threads**)
 - 3 write content of local TVars into global TVars
 - 4 remove the locks

For $calc \in \{SHF, CSHF\}$

Contextual Equivalence \sim_{calc}

$P_1 \sim_{calc} P_2$ iff for all contexts \mathbb{D} :

$$\mathbb{D}[P_1] \downarrow_{calc} \iff \mathbb{D}[P_2] \downarrow_{calc} \quad \wedge \quad \mathbb{D}[P_1] \Downarrow_{calc} \iff \mathbb{D}[P_2] \Downarrow_{calc}$$

where


- Process P is **successful** iff $P \equiv \mathbb{D}[\langle x \rangle u] \xleftarrow{\text{main}} \text{return } e]$

- **May-Convergence:**

$$P \downarrow_{calc} \text{ iff } \exists P' : P \xrightarrow{calc,*} P' \wedge P' \text{ is successful}$$

- **Should-Convergence:**

$$P \Downarrow_{calc} \text{ iff } \forall P' : P \xrightarrow{calc,*} P' \implies P' \downarrow_{calc}$$



Main Theorem

Convergence Equivalence: For any SHF -process P :

$$P \downarrow_{SHF} \iff \psi(P) \downarrow_{CSHF} \quad \text{and} \quad P \Downarrow_{SHF} \iff \psi(P) \Downarrow_{CSHF}$$

Adequacy: For all $P_1, P_2 \in SHF$:

$$\psi(P_1) \sim_{CSHF} \psi(P_2) \implies P_1 \sim_{SHF} P_2$$

- $CSHF$ is a **correct evaluator** for SHF
- Correct **program transformations** in $CSHF$ are also correct for SHF

Conclusion

- **Semantic correctness** of an STM-Haskell implementation
- using **contextual equivalence** with may- and should-convergence

Further work

- Transfer the result to **GHC's STM implementation**
- Develop smarter strategies for the transaction manager and prove their correctness
- Language extensions: **polymorphic** types, **exceptions**, . . .

Backup Slides

Conflict detection:

GHC STM: **thread** compares transaction log with content of TVars
restarts itself if a conflict occurred
(temporarily and before commit)

CSHF: the **committing thread restarts** conflicting **threads**

Pointer equality test:

GHC STM: required

CSHF : **not** required

Conflict requires:

GHC STM: **different** content

CSHF : **changed** content (not necessarily different)

- $P \downarrow_{SHF} \implies \psi(P) \downarrow_{CSHF}$:
map reductions $P \xrightarrow{SHF,*} P'$ to reductions $\psi(P) \xrightarrow{CSHF,*} \psi(P')$
- $\psi(P) \downarrow_{CSHF} \implies P \downarrow_{SHF}$:
 - reorder the sequence $\psi(P) \xrightarrow{CSHF,*} P'$, s.t. reductions are **grouped per transaction**
 - **remove non-committed** transactions
 - now the sequence can be mapped to a sequence $P \xrightarrow{SHF,*} P''$
- $P \downarrow_{SHF} \iff \psi(P) \downarrow_{CSHF}$:
 - similar, by showing equivalence of **may-divergence**:
 $P \uparrow_{SHF} \iff \psi(P) \uparrow_{CSHF}$
 - $P \uparrow = \neg(P \downarrow) = \exists Q : P \xrightarrow{*} Q \wedge \neg(Q \downarrow)$