

Compiler Correctness for Software Transactional Memory

Liyang HU

Foundations of Programming Group
School of Computer Science and IT
University of Nottingham

FoP Away Day, January 17th, 2007

Why Concurrency?

Limits of Technology

- Speed: 4GHz; plateaued over 2 years ago
- Power: 130W(!) from a die less than 15mm by 15mm
- Size: 65nm in 2006 – about 300 atoms across

Recent Trends

- Dual, even quad cores on a single package
- Multiprocessing has arrived for the *mass market*

Concurrent Programming (Is Hard!)

- Market leader: mutual exclusion
- Difficult to reason with

Example

Race Conditions

$deposit :: Account \rightarrow Integer \rightarrow IO ()$

$deposit\ account\ amount = do$

$\quad balance \leftarrow read\ account$

$\quad write\ account\ (balance + amount)$

Thread		<i>account</i> Balance
A	$balanceA \leftarrow read\ account$	<i>initial</i>
B	$balanceB \leftarrow read\ account$	<i>initial</i>
B	$write\ account\ (balanceB + amountB)$	$initial + amountB$
A	$write\ account\ (balanceA + amountA)$	$initial + amountA$

Example

Lack of Compositionality

deposit :: Account \rightarrow Integer \rightarrow IO ()

deposit account amount = do

lock account

balance \leftarrow read *account*

write *account* (*balance* + *amount*)

release account

transfer :: Account \rightarrow Account \rightarrow Integer \rightarrow IO ()

transfer from to amount = do

withdraw from amount

deposit to amount

Example

Lack of Compositionality (Solution?)

deposit :: Account → Integer → IO ()

deposit account amount = **do**

balance ← read account

 write account (*balance* + amount)

transfer :: Account → Account → Integer → IO ()

transfer from to amount = **do**

lock from; lock to

withdraw from amount

deposit to amount

release from; release to

Example

Deadlock

- Thread A: *transfer x y 100*
- Thread B: *transfer y x 200*

Thread		Account x	Account y	Action
A	<i>lock x</i>	free	free	acquires lock on x
B	<i>lock y</i>	held by A	free	acquires lock on y
B	<i>lock x</i>	held by A	held by B	waits for x...
A	<i>lock y</i>	held by A	held by B	waits for y...

Mutual Exclusion

Pitfalls

- Race conditions
- Priority inversion
- Deadlock
- Locking is often *advisory*

Drawbacks

- Correct code does not compose
- Overly conservative
- Granularity versus scalability

Transactional Model

What are Transactions?

- Arbitrary command sequence as an indivisible unit
- Declarative rather than descriptive
- Optimistic execution

Transactional Solution

```
work = do  
begin  
transfer a b 100  
commit
```


Advantages of Transactions

ACID Properties

- Atomicity: all or nothing
 - Fewer interleavings to consider
- Consistency: ensure invariants
 - System-enforced
- Isolation: no observable intermediate state
 - Guaranteed non-interference
- Durability: persistence through system failure
 - Simplifies error-handling
 - Not applicable for transactional *memory*

Optimistic Execution

Transactional Deposit

```

deposit :: Account → Integer → IO ()
deposit account amount = do
  begin
    balance ← read account
    -- another transaction commits, modifying account
    write account (balance + amount)
    commit -- fails
  
```

Failure and Retry

- DBMS tracks transaction *dependencies*
- External writes to *account* after initial *read* unacceptable
- Application can *retry* if aborted (not traditionally automatic)

Hardware Assistance

Atomic Instructions

- E.g. fetch-and-add, test-and-set
- Used to efficiently implement *mutual exclusion*

Avoiding Explicit Synchronisation

- Compare-and-Swap
 - CAS (a), b , c – if $(a) \equiv b$ then swap (a) with c
- Load-Linked / Store Conditional
 - Load-linked places watch on memory bus; begins 'transaction'
 - Access to watched location invalidates transaction
 - Store-conditional returns error code on failure
- Still not quite fully-fledged transactions

More Versatility?

Proposed Extensions

- Multi-word CAS
- Hardware Transactional Memory (Herlihy and Moss, 1993)
- Not available on a processor near you. . .

Software Transactional Memory

- Why wait for hardware? (Shavit and Touitou, 1995)
- Typical STM *libraries* difficult to use
- Language extension in Java (Harris and Fraser, 2003)

STM in Haskell

Composable Memory Transactions (Harris et al., 2005)

- Implemented in Glasgow Haskell Compiler
- Library and runtime system only; no language change

STM Haskell Primitives

```
instance Monad STM where { ... }
```

```
newTVar :: STM (TVar  $\alpha$ )
```

```
readTVar :: TVar  $\alpha$  → STM  $\alpha$ 
```

```
writeTVar :: TVar  $\alpha$  →  $\alpha$  → STM ()
```

```
retry    :: STM  $\alpha$ 
```

```
orElse   :: STM  $\alpha$  → STM  $\alpha$  → STM  $\alpha$ 
```

```
atomic   :: STM  $\alpha$  → IO  $\alpha$ 
```

Restricting Side-Effects

IO Actions

```
launchMissiles :: IO ()
```

```
atomic $ do
```

```
  launchMissiles  -- compile-time error: type mismatch
```

```
  ... retry ...
```

STM Monad

- Irreversible side-effects prohibited – the `IO` monad
- Can only read/write `TVars`
- But any *pure* code is allowed

Alternative Blocking

Try Again

- STM Haskell introduces the *retry* keyword
- Used where programs would block, or signal recoverable error

Composition

- *orElse* combines two transactions: *a 'orElse' b*
- Leftist: tries *a* first, returns if *a* returns
- If *a* calls *retry*, attempt *b*; one or the other succeeds

Code Flexibility

Blocking or Non-Blocking?

popBlocking :: TVar [Integer] → STM Integer

popBlocking ts = **do**

s ← *readTVar* ts

case *s* **of** [] → *retry*

(*x* : *xs*) → **do** *writeTVar* *xs*; *return* *x*

popNonblocking :: TVar [Integer] → STM (Maybe Integer)

popNonblocking ts = *liftM* Just (*popBlocking* ts)

'*orElse*' *return* Nothing

- Similarly turn non-blocking into blocking

Formal Semantics

Transition Rule for *atomic*

$$\frac{m \rightarrow^* \bar{n}}{\text{atomic } m \rightarrow \bar{n}}$$

The Need for a Low-Level Semantics

- Mixed big and small step semantics
 - No concurrent/optimistic execution of transactions
 - Doesn't use *logs*, as mentioned in the implementation
- Informal description of implementation
 - No attempt to relate to formal semantics
- How do we show *any* implementation correct?

Simplification of STM Haskell

Syntax

$E ::= \mathbb{Z} \mid E + E \mid \mathbf{rd} \text{ Name} \mid \mathbf{wr} \text{ Name } E \mid \mathbf{atomic} E$

Comparison with STM Haskell

STM Haskell

$(\gg=) \quad :: \text{STM } \alpha \rightarrow (\alpha \rightarrow \text{STM } \beta) \rightarrow \text{STM } \beta$

return $:: \alpha \rightarrow \text{STM } \alpha$

retry $:: \text{STM } \alpha$

orElse $:: \text{STM } \alpha \rightarrow \text{STM } \alpha \rightarrow \text{STM } \alpha$

readTVar $:: \text{TVar } \alpha \rightarrow \text{STM } \alpha$

writeTVar $:: \text{TVar } \alpha \rightarrow \alpha \rightarrow \text{STM } ()$

newTVar $:: \text{STM } (\text{TVar } \alpha)$

atomic $:: \text{STM } \alpha \rightarrow \text{IO } \alpha$

Model

$e + f$

$m \in \mathbb{Z}$

$\mathbf{rd} v$

$\mathbf{wr} v e$

$\mathbf{atomic} e$

Small-Step Semantics

$$\frac{\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle}{\langle e + f, \sigma \rangle \longrightarrow \langle e' + f, \sigma' \rangle} \quad (\text{ADDL})$$

$$\frac{\langle f, \sigma \rangle \longrightarrow \langle f', \sigma' \rangle}{\langle n + f, \sigma \rangle \longrightarrow \langle n + f', \sigma' \rangle} \quad (\text{ADDR})$$

$$\frac{}{\langle n + m, \sigma \rangle \longrightarrow \langle \overline{n + m}, \sigma \rangle} \quad (\text{ADDZ})$$

$$\frac{}{\langle \mathbf{rd} \ v, \sigma \rangle \longrightarrow \langle \sigma(v), \sigma \rangle} \quad (\text{READ})$$

$$\frac{\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle}{\langle \mathbf{wr} \ v \ e, \sigma \rangle \longrightarrow \langle \mathbf{wr} \ v \ e', \sigma' \rangle} \quad (\text{WRITEE})$$

$$\frac{\langle e, \sigma \rangle \longrightarrow^* \langle n, \sigma' \rangle}{\langle \mathbf{atomic} \ e, \sigma \rangle \longrightarrow \langle n, \sigma' \rangle} \quad (\text{ATOMIC})$$

$$\frac{}{\langle \mathbf{wr} \ v \ \bar{n}, \sigma \rangle \longrightarrow \langle \sigma(v), \sigma[v \mapsto n] \rangle} \quad (\text{WRITEZ})$$

Concurrent Evaluation

Expression Soup

$$P ::= E \mid P \parallel P$$

$$\frac{\langle e, \sigma \rangle \longrightarrow \langle e', \sigma' \rangle}{\langle e, \sigma \rangle \longrightarrow \langle e, \sigma' \rangle} \quad (\text{SEQ})$$

$$\frac{\langle p, \sigma \rangle \longrightarrow \langle p', \sigma' \rangle}{\langle p \parallel q, \sigma \rangle \longrightarrow \langle p' \parallel q, \sigma' \rangle} \quad (\text{PARL})$$

$$\frac{\langle q, \sigma \rangle \longrightarrow \langle q', \sigma' \rangle}{\langle p \parallel q, \sigma \rangle \longrightarrow \langle p \parallel q', \sigma' \rangle} \quad (\text{PARR})$$

Example

- `rd "x" + rd "x" ∥ wr "x" 1` — yields 0, 1 or 2
- `atomic (rd "x" + rd "x") ∥ wr "x" 1` — yields only 0 or 2

Virtual Machine

Instruction Set

Instruction ::=	PUSH \mathbb{Z}		ADD	-- stack machine
	LOAD Name		SWAP Name	-- shared store
	BEGIN		COMMIT	-- transactions

- Typical stack machine with a shared store
- LOAD and SWAP are transaction-local if one is active
- BEGIN marks the start of a transaction
- COMMIT marks the end; retries on failure

Implementation?

- Easiest: stop-the-world; no interleaving of transactions

Logs and Transaction Frames

Goals

- 1 Isolate changes to global state
- 2 Re-run transaction on abort

Transaction Frame

We need to record:

- 1 for each variable accessed,
 - its original value – to check for conflicting commits; and
 - value of writes to it – subsequent reads return this value
- 2 the transaction's starting address – to re-run if commit fails
- 3 and strictly speaking, the stack too...

Each frame is a pair

$\langle ip, rw \rangle \in \text{TransactionFrame} \equiv \text{Instruction}^* \times (\text{Name} \rightarrow \mathbb{Z} \times \mathbb{Z})$

Concurrent Execution

Threads

$$\langle ip, sp, tp \rangle \in \text{Thread} \equiv \text{Instruction}^* \times \mathbb{Z}^* \times \text{TransactionFrame}^*$$

Thread Soup

Program ::= Thread
 | Program \square Program

- Rules (SEQ), (PARL) and (PARR) will suffice
- Threads execute paired with a shared store

Compiler

E to Instruction*

$$\text{compE} \in E \rightarrow \text{Instruction}^* \rightarrow \text{Instruction}^*$$

$$\text{compE } \bar{n} \quad c = \text{PUSH } n : c$$

$$\text{compE } (e + f) \quad c = \text{compE } e \ (\text{compE } f \ (\text{ADD} : c))$$

$$\text{compE } (\text{rd } v) \quad c = \text{LOAD } v : c$$

$$\text{compE } (\text{wr } v \ e) \quad c = \text{compE } e \ (\text{SWAP } v : c)$$

$$\text{compE } (\text{atomic } e) \quad c = \text{BEGIN} : \text{compE } e \ (\text{COMMIT} : c)$$

P to Program

$$\text{compP} \in P \rightarrow \text{Program}$$

$$\text{compP } e \quad = \langle |\text{compE}e|, |e|, |e| \rangle$$

$$\text{compP } (p \parallel q) = \text{compP } p \parallel \text{compP } q$$

Correctness

Sequential

$\forall e \in E, \sigma \in \text{Name} \rightarrow \mathbb{Z}, n \in \mathbb{Z}.$

$$\langle e, \sigma \rangle \longrightarrow^* \langle \bar{n}, \sigma' \rangle$$

iff

$$\langle \langle \text{comp}E e [], [], [] \rangle, \sigma \rangle \longrightarrow^* \langle \langle [], [n], [] \rangle, \sigma' \rangle$$

Concurrent

$\forall p \in P, \sigma \in \text{Name} \rightarrow \mathbb{Z}, ns \in P.$

$$\langle p, \sigma \rangle \longrightarrow^* \langle ns, \sigma' \rangle$$

iff

$$\langle \text{comp}P p, \sigma \rangle \longrightarrow^* \langle rs, \sigma' \rangle$$

- $ns \in P$ contains only integer expressions of the form \bar{n}
- $rs \in \text{Program}$ structurally identical to ns but with $\bar{n} \mapsto \langle [], [n], [] \rangle$

Model Verification

Implementation

- Small-step semantics, compiler and VM in Haskell
- Can express compiler correctness as following function:

$propCC :: P \rightarrow Bool$

$propCC\ p = (result \circ run)\ (p, \sigma_0) \equiv (result \circ run)\ (compP\ p, \sigma_0)$

QuickCheck

- Generates random input, attempts to falsify proposition:
 - > `quickCheck propCC`
 - OK, passed 100 tests.
 - >
- Inspires confidence that a formal proof is possible...

Interference and Serialisability

Questions

- What kind of interference can we allow?
- How do we serialise transactions? When do they 'happen'?

Interfering Transactions

Thread				TVars	
A	B	C	D	x	y
<div style="border: 1px solid black; padding: 5px;"> rd $x \mapsto 0$ rd $y \mapsto 1$... <i>commit?</i> </div>	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> wr $x \ 1$ </div>	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> wr $y \ (rd \ x \mapsto 1)$ </div>	<div style="border: 1px solid black; padding: 5px; width: fit-content; margin: 10px auto;"> wr $x \ 0$ </div>	0	0
				1	0
				1	1
				1	1
				0	1
				0	1

Optimistic Speculation

Answers

- Permitted interference?
 - On initial access, bet on variable's final pre-commit value
 - Allow *any* changes, provided original value restored
 - If so, the transaction commits successfully
- At what point does a transaction take place?
 - Certainly not when the transaction begins
 - Pre-commit, x and y matches what thread A initially read
 - Hence, can collapse down to the successful commit point

Read / Write Reordering

- Reads happen immediately
- Writes buffered until commit time
- Commit behaves almost like MCAS

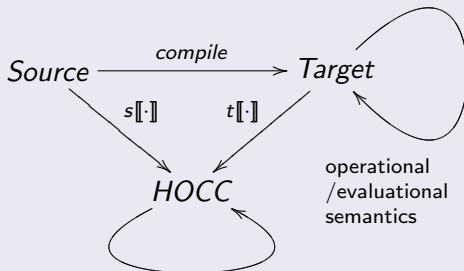
On Equality

Equality Strengths

- Value, or structural
 - Fast for primitive values, bad for lazy thunks
- Pointer
 - Efficient for unevaluated thunks and primitive values
 - Can't replace value by a copy of the same
- Version
 - Considers writes without regard to actual values involved
 - By pairing values with an incrementing version number
 - Or by a watch on the memory location, c.f. LL/SC
- State
 - All changes to shared state undesirable
- World
 - All interleaving undesirable

Existing Methodology

Compiler Correctness for Parallel Languages (Wand, 1995)



- Compiler correct if $s[p]$ bisimilar to $t[\text{compile } p]$
- Target operational semantics adequate relative to HOCC

Something Simpler?

Aim and Overview

- Avoid so many layers of translation; too much room for error
- Give source/target languages small-step/operational semantics
- Augment semantics with labelled transition system
- Direct bisimulation between the two semantics

Expressions and Evaluation

Expressions

$$E ::= \mathbb{Z} \mid E + E$$

- Addition supplemented with a (ZAP) rule
- Simple form of non-determinacy
- Left-biased evaluation

Labelled Transition System

$$\text{Action} ::= \mathbb{Z} + \mathbb{Z} \mid \mathbb{Z} \downarrow \mathbb{Z}$$

$$\text{Label} ::= \text{Action} \mid \tau$$

$$\dot{\rightarrow} \subseteq E \times \text{Label} \times E$$

Evaluation

Reduction Rules

$$\frac{}{\bar{n} + \bar{m} \xrightarrow{n+m} \overline{n+m}} \quad (\text{ADD})$$

$$\frac{}{\bar{n} + \bar{m} \xrightarrow{n \neq m} 0} \quad (\text{ZAP})$$

$$\frac{e \xrightarrow{\alpha} e'}{e + f \xrightarrow{\alpha} e' + f} \quad (\text{ADDL})$$

$$\frac{f \xrightarrow{\alpha} f'}{\bar{n} + f \xrightarrow{\alpha} \bar{n} + f'} \quad (\text{ADDR})$$

Choice of Action

- Differentiate base case reductions in source language
- Two symbols are enough but...
- Conceivably, a broken compiler could keep structure intact
- Include operands to ensure the same values are computed

Compiler

Virtual Machine

$$I ::= \text{PUSH } \mathbb{Z} \mid \text{ADD}$$

$$M = I^* \times \mathbb{Z}^*$$

$$\dot{\rightarrow} \subseteq M \times \text{Label} \times M$$

Compiler

$$\text{compile} :: E \rightarrow I^* \rightarrow I^*$$

$$\text{compile } \bar{n} \quad \text{is} = \text{PUSH } n : \text{is}$$

$$\text{compile } (x + y) \text{ is} = \text{compile } x \text{ is}'$$

where $\text{is}' = \text{compile } y \text{ (ADD : is)}$

Execution

Virtual Machine Transitions

$$\langle \text{PUSH } n : \iota\sigma, \sigma \rangle \xrightarrow{\tau} \langle \iota\sigma, n : \sigma \rangle \quad (\text{PUSH})$$

$$\langle \text{ADD} : \iota\sigma, m : n : \sigma \rangle \xrightarrow{n+m} \langle \iota\sigma, n + m : \sigma \rangle \quad (\text{ADD})$$

$$\langle \text{ADD} : \iota\sigma, m : n : \sigma \rangle \xrightarrow{n \not\leq m} \langle \iota\sigma, 0 : \sigma \rangle \quad (\text{ZAP})$$

- Similar non-deterministic semantics, c.f. (ADD) and (ZAP)

Mixed Bisimulation

Motivation

- Can express correctness as $\langle \text{compile } x \llbracket \cdot \rrbracket, \llbracket \cdot \rrbracket \rangle \approx x$
 - At every reduction step, anything LHS can do, RHS can follow
- Proof for something like this: structural induction on e ?
- Need to generalise on stack, instruction continuation...
- Introduce expression *contexts*, $c[\llbracket \cdot \rrbracket]$?
- Can certainly relate stack and continuation to context
 - But proof turns very messy; this is a simple language!

Combined Machine – Existing Technology!

$$C \equiv (E + 1) \times M$$

$$\rightarrow \subseteq C \times \text{Label} \times C$$

Combined Semantics

Transition Rules

$$\frac{x \xrightarrow{\alpha} x'}{\langle x, \text{ls}, \sigma \rangle \xrightarrow{\alpha} \langle x', \text{ls}, \sigma \rangle} \quad (\text{EVAL})$$

$$\frac{}{\langle \bar{n}, \text{ls}, \sigma \rangle \xrightarrow{\tau} \langle \bullet, \text{ls}, n : \sigma \rangle} \quad (\text{SWITCH})$$

$$\frac{\langle \text{ls}, \sigma \rangle \xrightarrow{\alpha} \langle \text{ls}', \sigma' \rangle}{\langle \bullet, \text{ls}, \sigma \rangle \xrightarrow{\alpha} \langle \bullet, \text{ls}', \sigma' \rangle} \quad (\text{EXEC})$$

Weak Simulation

Definition

A non-empty relation $\mathfrak{R} \subseteq C \times C$ is a *weak simulation* iff for all $c \mathfrak{R} d$,

$$c \xrightarrow{\alpha} c' \quad \text{implies} \quad \exists d'. d \xrightarrow{\alpha} d' \wedge c' \mathfrak{R} d'$$

- There exists a maximal \mathfrak{R} : we name it \succcurlyeq
- $c \succcurlyeq d$ and $c \preccurlyeq d$ iff $c \approx d$

Lemma (Eliding τ)

If $c \xrightarrow{\tau} c'$ is the only possible transition by c , then:

$$\frac{c \xrightarrow{\tau} c'}{c \preccurlyeq c'} \quad \text{and} \quad \frac{c \xrightarrow{\tau} c'}{c \succcurlyeq c'}, \quad \text{or} \quad \frac{c \xrightarrow{\tau} c'}{c \approx c'}$$

Compiler Correctness

Theorem 1 (Soundness)

$$\langle x, \text{is}, \sigma \rangle \succcurlyeq \langle \bullet, \text{compile } x \text{ is}, \sigma \rangle$$

Everything program does permitted by expression semantics

Proof Overview

- In this case, soundness and completeness proofs are identical
 - Recover separate proofs by replacing \approx with \preccurlyeq or \succcurlyeq
- Completeness may not always be possible or even required
- Corollary (Correctness): $\langle x, [], [] \rangle \approx \langle \bullet, \text{compile } x [], [] \rangle$
- Selected highlights follow...
 - For full details, see my first year transfer dissertation

Compiler Correctness

Theorem 2 (Completeness)

$$\langle x, \text{is}, \sigma \rangle \preceq \langle \bullet, \text{compile } x \text{ is}, \sigma \rangle$$

Program does everything permitted by expression semantics

Proof Overview

- In this case, soundness and completeness proofs are identical
 - Recover separate proofs by replacing \approx with \preceq or \succeq
- Completeness may not always be possible or even required
- Corollary (Correctness): $\langle x, [], [] \rangle \approx \langle \bullet, \text{compile } x [], [] \rangle$
- Selected highlights follow...
 - For full details, see my first year transfer dissertation

Compiler Correctness

Theorem 3 (Bisimulation)

$$\langle x, \text{is}, \sigma \rangle \approx \langle \bullet, \text{compile } x \text{ is}, \sigma \rangle$$

Program is a bisimulation of expression semantics

Proof Overview

- In this case, soundness and completeness proofs are identical
 - Recover separate proofs by replacing \approx with \preceq or \succeq
- Completeness may not always be possible or even required
- Corollary (Correctness): $\langle x, [], [] \rangle \approx \langle \bullet, \text{compile } x [], [] \rangle$
- Selected highlights follow...
 - For full details, see my first year transfer dissertation

Theorem 3: $\langle x, \mathcal{I}, \sigma \rangle \approx \langle \bullet, \text{compile } x \mathcal{I}, \sigma \rangle$

Inductive Case: $x \equiv y + z$

Have induction hypothesis for y :

$$\forall \mathcal{I}', \sigma'. \langle y, \mathcal{I}', \sigma' \rangle \approx \langle \bullet, \text{compile } y \mathcal{I}', \sigma' \rangle$$

and also for z . Then:

$$\begin{aligned} & \langle \bullet, \text{compile } (y + z) \mathcal{I}, \sigma \rangle \\ \equiv & \quad \{ \text{definition of } \text{compile} \} \\ & \langle \bullet, \text{compile } y (\text{compile } z (\text{ADD} : \mathcal{I})), \sigma \rangle \\ \approx & \quad \{ \text{induction hypothesis for } y \} \\ & \langle y, \text{compile } z (\text{ADD} : \mathcal{I}), \sigma \rangle \\ \approx & \quad \{ \text{by lemma 4, given induction hypothesis for } z \} \\ & \langle y + z, \mathcal{I}, \sigma \rangle \end{aligned}$$

Additional Lemmas

Lemma 4 (Evaluate Left)

Given $\langle \bullet, \text{compile } z \text{ is}', \sigma' \rangle \approx \langle z, \text{is}', \sigma' \rangle,$

$$\langle y, \text{compile } z \text{ (ADD : is)}, \sigma \rangle \approx \langle y + z, \text{is}, \sigma \rangle$$

Proof – case $y \neq \bar{m}$

LHS:

$$\frac{y \xrightarrow{\alpha} y'}{\langle y, \text{compile } z \text{ (ADD is)}, \sigma \rangle} \text{ (EVAL)}$$

$$\xrightarrow{\alpha} \langle y', \text{compile } z \text{ (ADD : is)}, \sigma \rangle$$

RHS:

$$\frac{\frac{y \xrightarrow{\alpha} y'}{y + z \xrightarrow{\alpha} y' + z} \text{ (ADDL)}}{\langle y + z, \text{is}, \sigma \rangle \xrightarrow{\alpha} \langle y' + z, \text{is}, \sigma \rangle} \text{ (EVAL)}$$

Additional Lemmas

Lemma 5 (Evaluate Right)

$$\langle z, \text{ADD} : \iota s, \bar{m} : \sigma \rangle \approx \langle \bar{m} + z, \iota s, \sigma \rangle$$

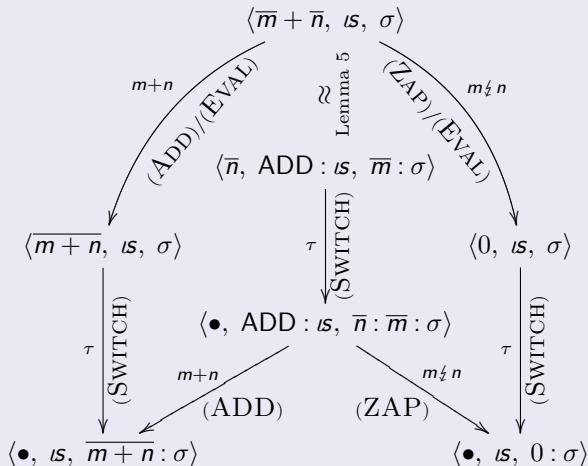
- The $z \neq \bar{n}$ case proceeds as lemma 4

Proof Method

- Uses simple *equational reasoning* and logic
- No need to consider *sets* of machine states / expressions
- Where there is non-determinism, we can chase diagrams
 - Weak bisimulation: traces $\alpha\tau$ and $\tau\alpha$ are equivalent

Chasing Diagrams

Proof of lemma 5 – case $z \equiv \bar{n}$



Conclusion

Future Work

- Extension of language with parallelism
- Exceptions and interrupts
- Proof of STM model
- Richer transactional memory constructs?
 - Forking within transactions
 - Compensating transactions
 - Data invariants