Transactional Futures Parallelism in Software Transactions

Janwillem Swalens





Context: concurrency





Concurrency is **useful** but **difficult**

Concurrency models provide safety guarantees

Three concurrency models

- Futures: deterministic
- **Transactions:** non-deterministic, shared memory
- Actors: non-deterministic, message passing

Separately they work fine, but combining them leads to problems.

Three concurrency models

- **Futures**: deterministic
- **Transactions**: non-deterministic, shared memory
- Actors: non-deterministic, message passing

Separately they work fine, but combining them leads to problems.

Futures for parallelism

(fork e) returns f

(join f) returns result of e

Futures for parallelism

(fork e) returns f

(join f) returns result of e



```
Transactions for
          shared memory
                            serializability
(ref v)
(atomic e)
(deref r)
(ref-set r v)
(def checking (ref 100))
(def savings (ref 500))
(fork
 (atomic
   (ref-set checking (- (deref checking) 10))
   (ref-set savings (+ (deref savings) 10)))
(fork
 (atomic
   (println "You own €" (+ (deref checking)
                          (deref savings)))))
```

Nesting futures & transactions



			2d		2s
	1s				
		4d			
				1d	
3s					
		3d			4s

			2d	2	2s
	1s	1	1	1	
		4d		1	
		4		1d	
3s		4	4	4	4
3	3	3d			4s















Labyrinth has limited speed-up

Problems when creating threads in a transaction

 Threads in transaction do not share context (Clojure, ScalaSTM)
 ⇒ no access to transactional state or ⇒ serializability violated

atomically

do { **forkIO** ... }

Transactional Tasks

Parallelism in transaction

⇒ Transactional task = thread created in transaction

Task can access transactional variables

⇒ Task adopts encapsulating transactional context

Isolation between tasks

⇒ Tasks work on conceptual copy

Serializability

⇒ All tasks should join before transaction commits On conflict, *all* tasks abort

Task = snapshot + store

(atomic (ref-set ... 1)

Each transactional task contains:

snapshot σ transactional state on creation **local store** τ local modifications

 $\begin{array}{l} \mathsf{T}_{\mathsf{x}} \cup \langle f_{p}, \overline{\sigma}, \tau, \mathsf{F}_{\mathsf{s}}, & \mathsf{F}_{\mathsf{j}}, \mathcal{E}[\texttt{fork } e] \rangle \\ \Rightarrow_{\mathsf{tf}} \mathsf{T}_{\mathsf{x}} \cup \langle f_{p}, \sigma, \tau, \mathsf{F}_{\mathsf{s}} \cup \{f'\}, \mathsf{F}_{\mathsf{j}}, \mathcal{E}[f_{c}] & \rangle \cup \langle f_{c}, \sigma :: \tau, \varnothing, \varnothing, \mathsf{F}_{\mathsf{j}}, e \rangle \\ & \text{with } f_{c} \text{ fresh} \end{array}$

fork creates isolated task

```
(atomic
  (ref-set ... 1)
  (fork
      (ref-set ... 2))
  (ref-set ... 2)
```

Each transactional task contains:

snapshot σ : transactional state on creation **local store** τ : local modifications

 $T_{x} \cup \langle f_{p}, \sigma, \tau, F_{s}, F_{j}, \mathcal{E}[\text{fork } e] \rangle$ $\Rightarrow_{\text{tf}} T_{x} \cup \langle f_{p}, \sigma, \tau, F_{s} \cup \{f'\}, F_{j}, \mathcal{E}[f_{c}] \rangle \cup \langle f_{c}, \sigma :: \tau \otimes \emptyset, F_{j}, e \rangle$ with f_{c} fresh

join merges changes

(atomic

(join child))

merge local store τ' of child into parent

 $T_{x} \cup \langle f_{p}, \sigma, \tau \rangle F_{s}, F_{j}, \qquad \mathcal{E}[\operatorname{join} f_{e}] \rangle \cup \langle f_{c}, \sigma', \tau', F_{s}', F_{j}', v \rangle$ $\Rightarrow_{tf} T_{x} \cup \langle f_{p}, \sigma, \tau :: \tau') F_{s}', F_{j} \cup F_{j}' \cup \{f_{c}\}, \mathcal{E}[v] \qquad \rangle \cup \langle f_{c}, \sigma', \tau', F_{s}', F_{j}', v \rangle$ $\text{if } f_{c} \notin F_{j} \text{ and } F_{s}' \subseteq F_{j}'$

Conflict resolution function:

(ref 0 resolve)

(defn resolve [o p c] (min p c)) (defn resolve [o p c] c) (defn resolve [o p c] (error "merge conflict"))

Properties of transactional tasks

- In-transaction parallelism possible
- Serializability of transactions
- Coordination of tasks: all or none
- In-transaction determinacy

Evaluation: Labyrinth

Evaluation: Labyrinth

Evaluation: Bayes

Evaluation: Bayes

Insights from experiments

- Labyrinth: parallelize search algorithm
 ⇒ fewer & cheaper conflicts
- Bayes: more fine-grained parallelism
 ⇒ better exploit hardware
- Low developer effort (re-use existing concepts)
- Suitable for applications with long transactions

Implementation

Fork of Clojure

https://github.com/jswalens/transactional-futures/

http://soft.vub.ac.be/~jswalens/ecoop-2016-artifact/

Summary

Parallelism in a transaction is useful for programs with **long transactions** But currently:

- / not allowed (Haskell)
- **∖ not serializable** (Clojure, Scala)

Idea: transactional tasks

- safe access to encapsulating transaction
- serializable, coordinated, determinate

Benefits:

- finer-grained parallelism \Rightarrow speed-up
- low developer effort

More details:

- J. Swalens et al. "Transactional Tasks: Parallelism in Software Transactions", ECOOP 2016
- https://github.com/jswalens/transactional-futures/

Maximal number of

Transactional Tasks vs. Nested Parallel Transactions

Guarantees in transaction

- NPT: (atomic (fork ...)) → race conditions possible
- NPT: (atomic (fork (atomic ...))) → serializable, last writer wins (not deterministic)
- TT: conflict resolution → *in-transaction determinacy* (but may need to define resolution function)

<u>Performance</u>

- NPT: roll back and retry subtransaction
- TT: resolve conflict
- → different performance characteristics depending on application (chance of conflicts between threads in tx)

More fine-tuned conflict resolution (e.g. minimum for Labyrinth)

STAMP

Application	Instructions	Time	
	/tx (mean)	in tx	
labyrinth	219,571	100%	
bayes	$60,\!584$	83% 🔴	
yada	9,795	100% $ullet$	
vacation-high	3,223 ●	86% $ullet$	
genome	1,717 ●	97% $ullet$	
intruder	$330 \bigcirc$	33% 🌓	
kmeans-high	$117~{\odot}$	7% \bigcirc	
ssca2	$50 \ \bigcirc$	17% \bigcirc	

Coarse-grained parallelism

between parts of the application

Transactions

- Conflicts span multiple variables
- Difficult to define conflict resolution functions
- Chance of conflicts depends on application
- ⇒ resolve high-level conflicts using serializability

Fine-grained parallelism within a part of the application

Transactional tasks

- Conflict affects single variable
- Define conflict resolution function based on algorithm
- Conflicts likely, so rollback bad for performance
- ⇒ resolve low-level conflicts using conflict resolution functions

Implementation details

(a) Code example.

```
(atomic

<sup>1</sup>(ref-set gray A)

<sup>2</sup>(fork <sup>3</sup>(ref-set blue B)

<sup>4</sup>(fork <sup>5</sup>(ref-set green C))

<sup>6</sup>(ref-set red D)

...)

<sup>7</sup>(ref-set purple E)

...)
```

(b) Data after step 3.

(c) Data after step 7.

