

Transactional Tasks: Parallelism in Software Transactions

Janwillem Swalens

Futures

(**fork** e) returns f

(**join** f) returns result of e

```
(defn fib [n]
  (if (< n 2)
      n
      (let [a (fib (- n 1))
            b (fib (- n 2))]
          (+ a b))))
```

Futures

(**fork** e) returns f

(**join** f) returns result of e

```
(defn fib [n]
  (if (< n 2)
      n
      (let [a (fork (fib (- n 1)))
            b (fork (fib (- n 2)))]
          (+ (join a) (join b)))))
```

Transactions

```
(ref v)
(atomic e)
(deref r) or @r
(ref-set r v)
```

serializability

```
(def checking (ref 100))
(def savings (ref 500))
(fork
  (atomic
    (ref-set checking (- @checking 10))
    (ref-set savings (+ @savings 10))))
(fork
  (atomic
    (println "You own €" (+ @checking @savings))))
```

Example: Labyrinth

			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

Example: Labyrinth

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

		2	3		
	s	1	2	3	
2	1	2	3		
				d	

		2	3	4	5
	s	1	2	3	
2	1	2	3	4	5
3		3	4	d	
4		4	5		
5		5			

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

; for each (src, dst) pair:

```
(let [local-grid (copy grid)]
  (expand src dst local-grid)
  (add-path grid
    (traceback local-grid dst)))
```

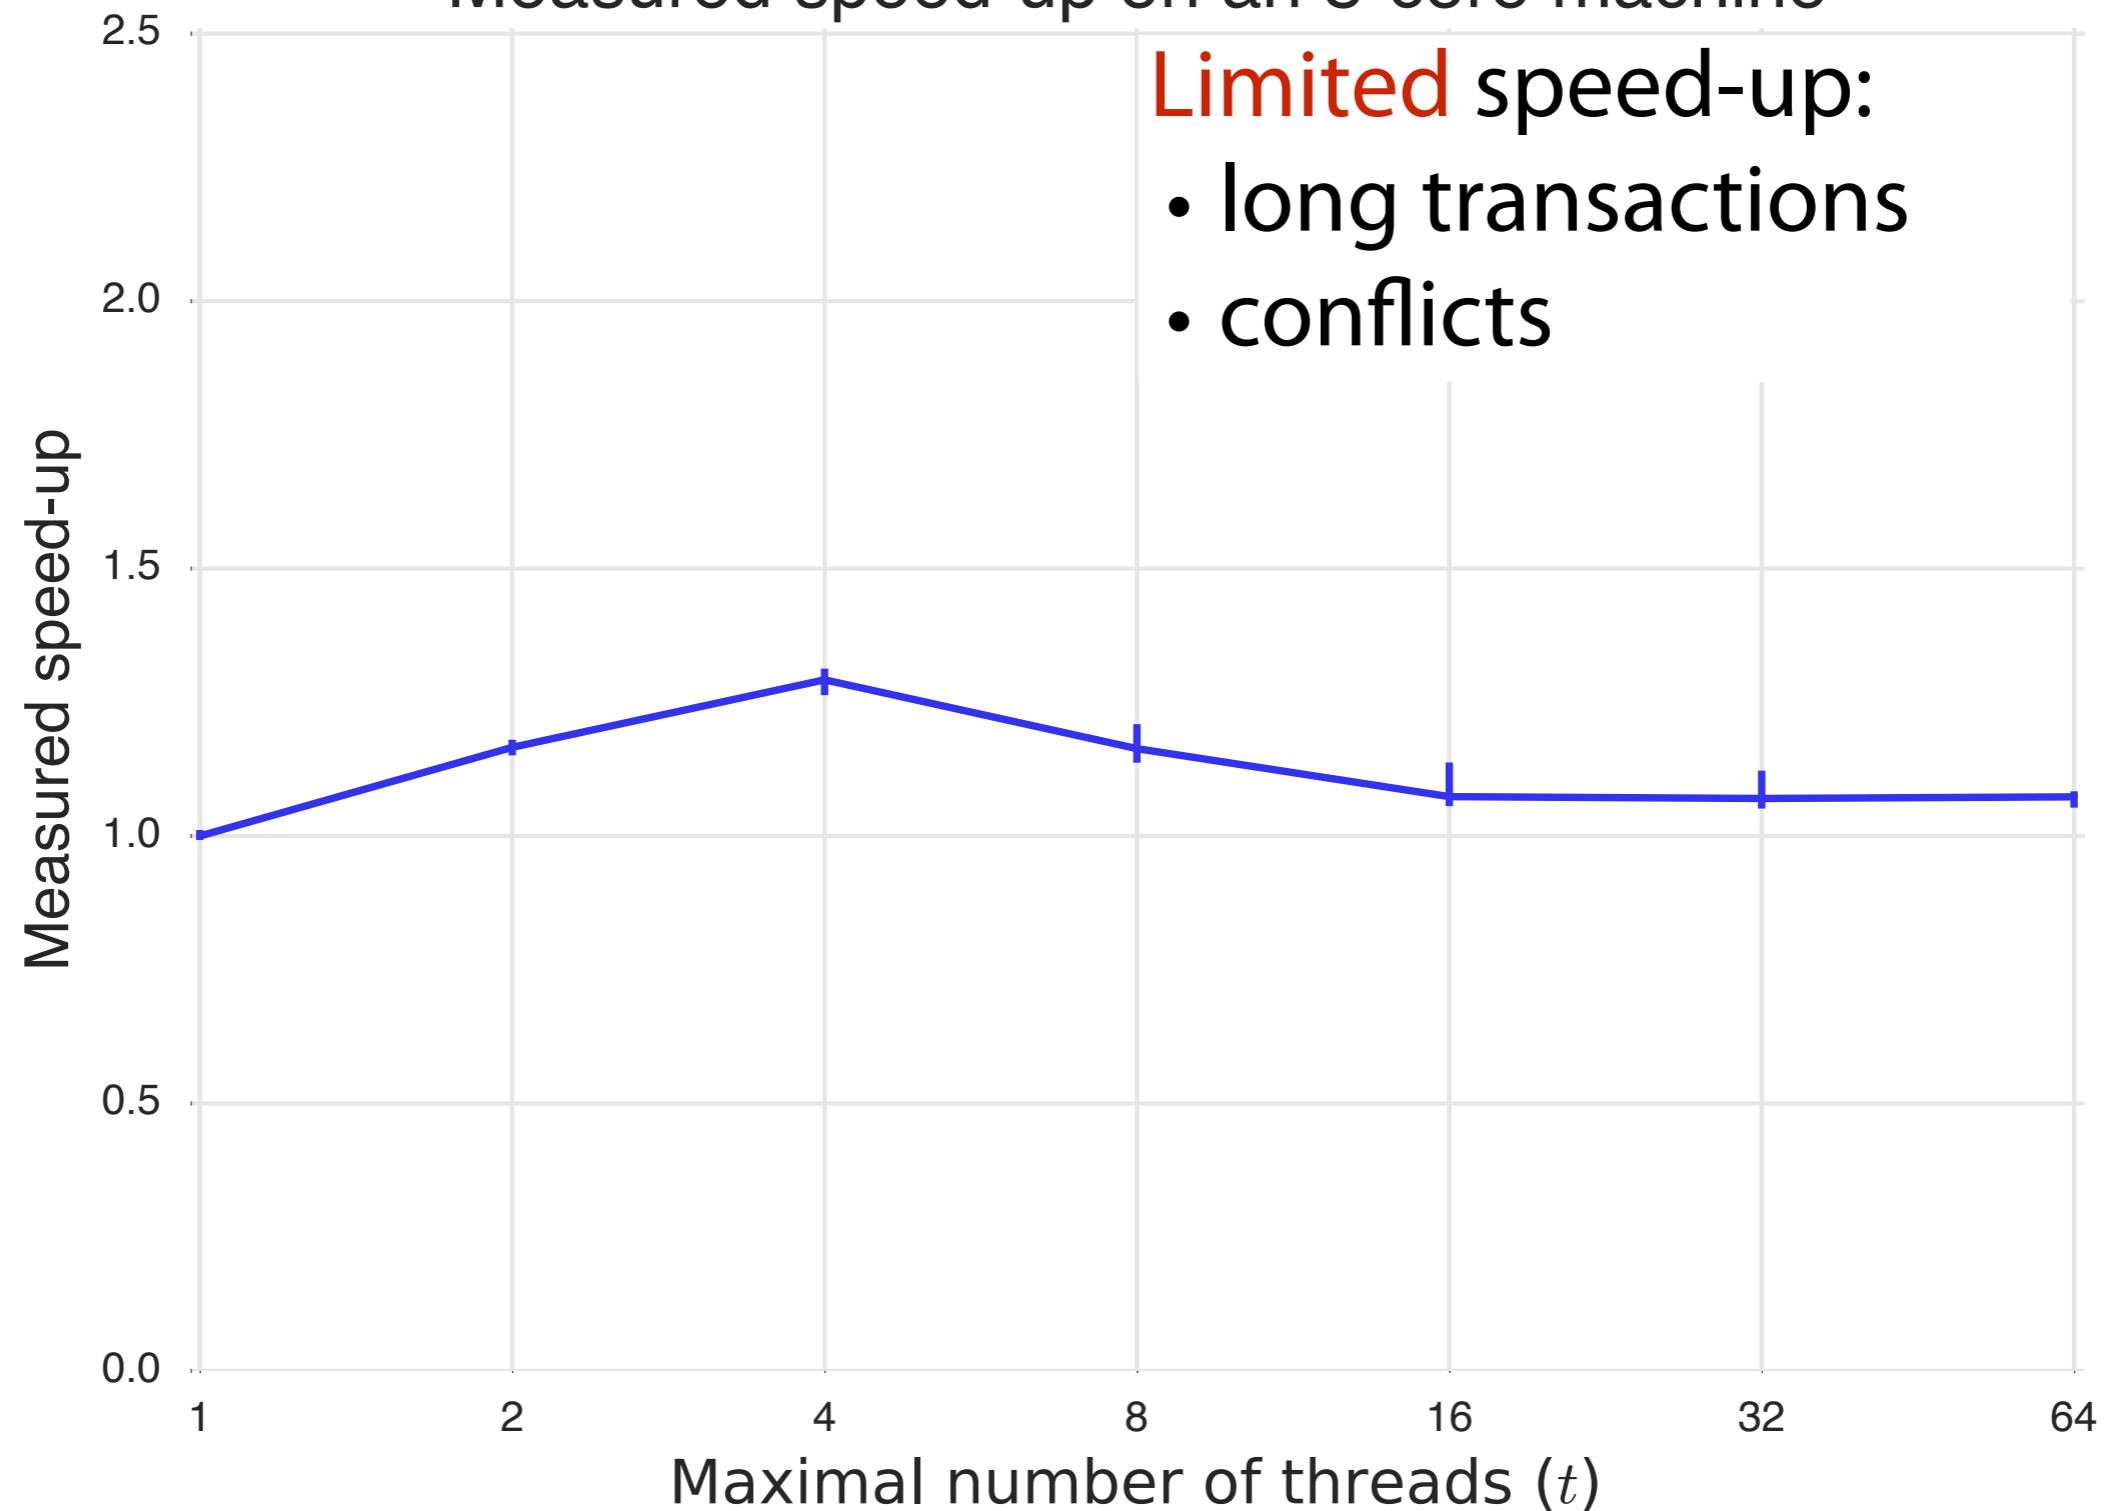
```
(defn expand [src dst
  (loop [q]
    (if (empty? q)
        false
        true
        (rest q)
        (expand-point (first q)
```

breadth-first search

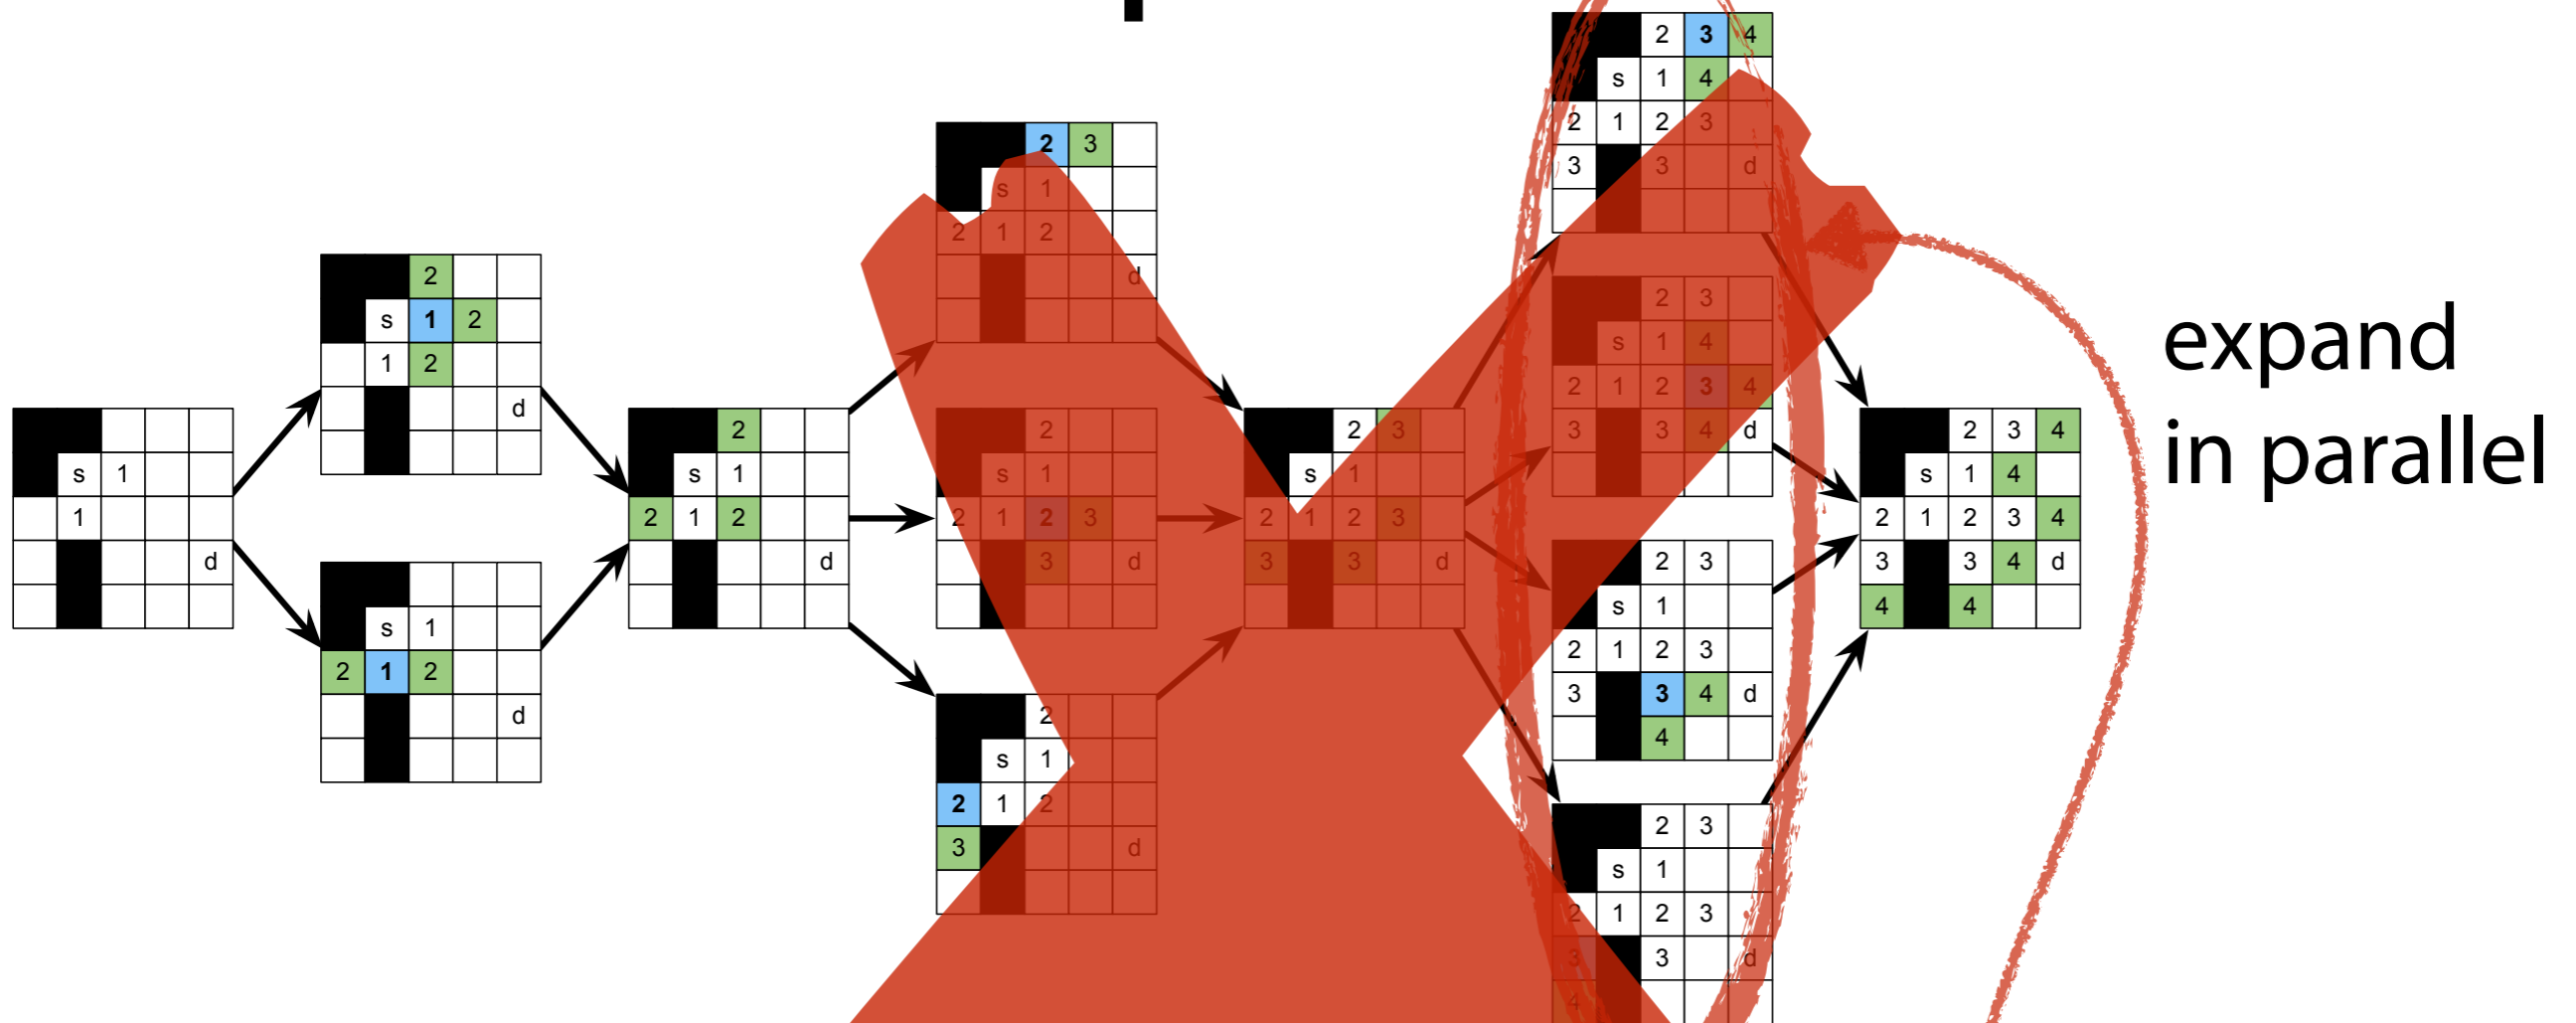
```
(defn expand-point [pt grid]
  (atomic
    (let [neighbors ...]
      (for [n neighbors]
        (ref-set n ...))
      neighbors)))
```

Labyrinth speed-up

Measured speed-up on an 8-core machine



Solution: parallel search



```

; for each (src, dst) pair:
(atomic
  (let [local-grid (copy grid)]
    (expand src dst local-grid)
    (add-path grid
      (traceback local-grid dst))))

```

```

(defn expand [src dst
  (loop [q #{src}]
    (if (empty? q)
        false
        (if (contains? q dst)
            true
            (recur
              (expand-step q

```

```

(defn expand-step [q grid]
  (reduce union #{}
    (pmap
      (fn [p] (expand-point p grid))
      q)))

```

parallel breadth-first search

Does not work!

Problems

- Threads in tx prohibited (Haskell)

⇒ **parallelism limited**

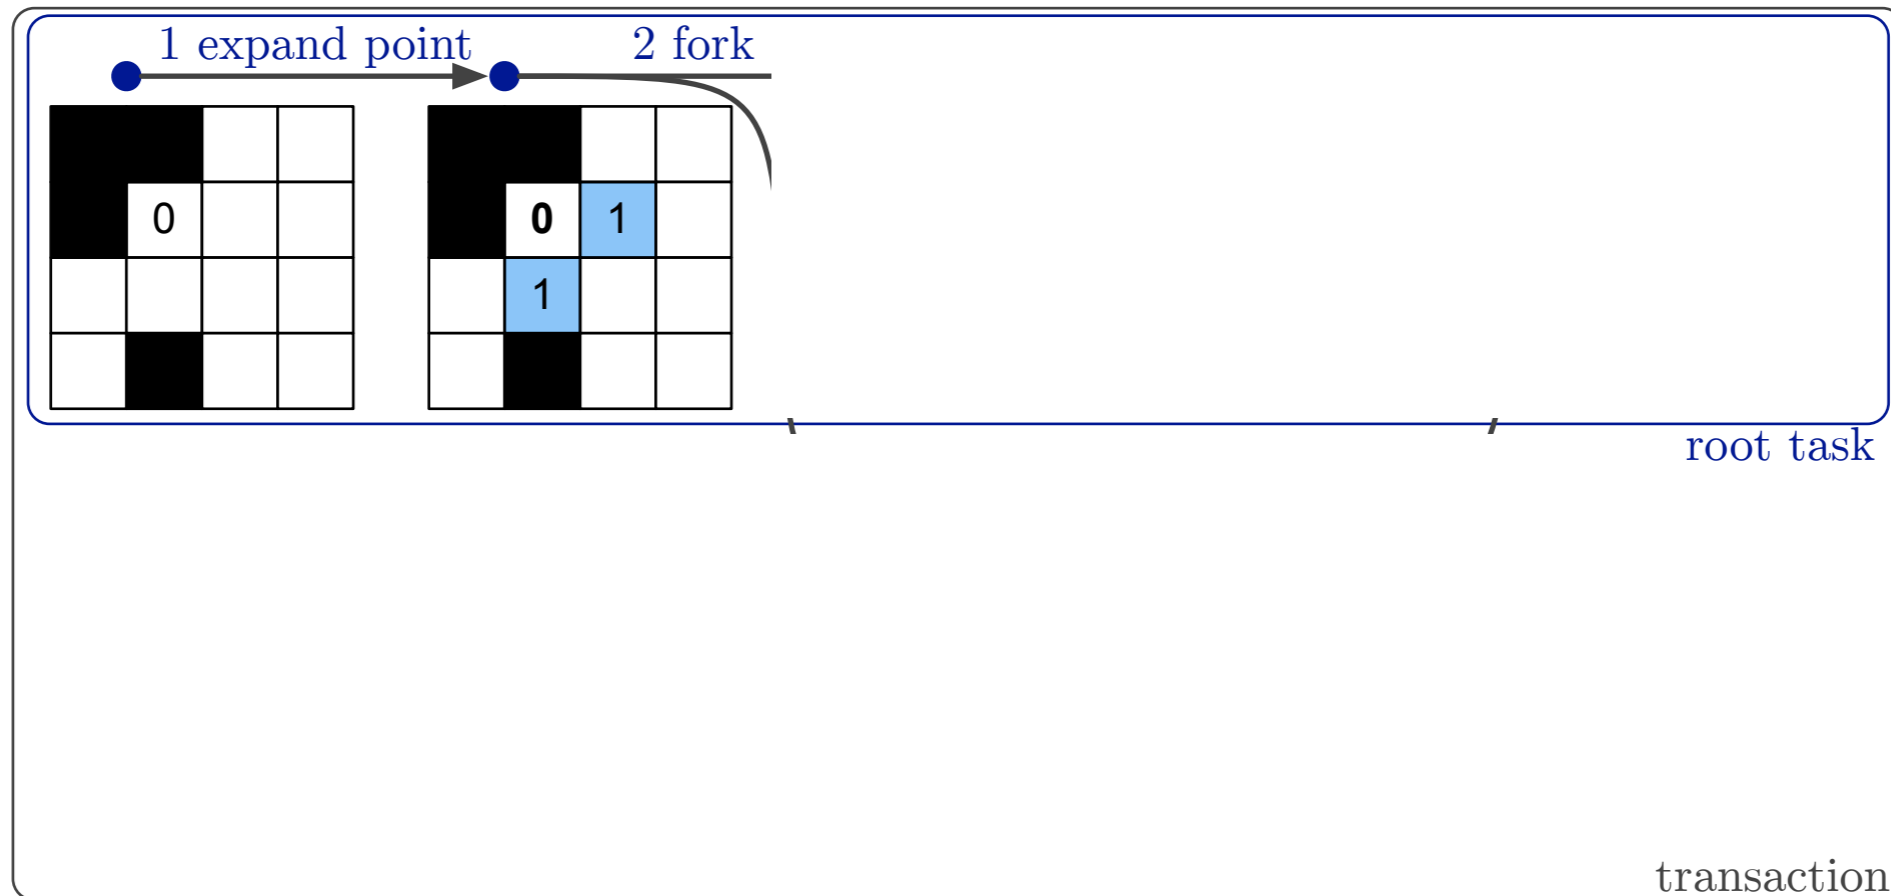
```
atomically $  
do { forkIO ... }
```

- Threads in tx allowed but don't share context (Clojure, Scala)

⇒ **serializability violated**

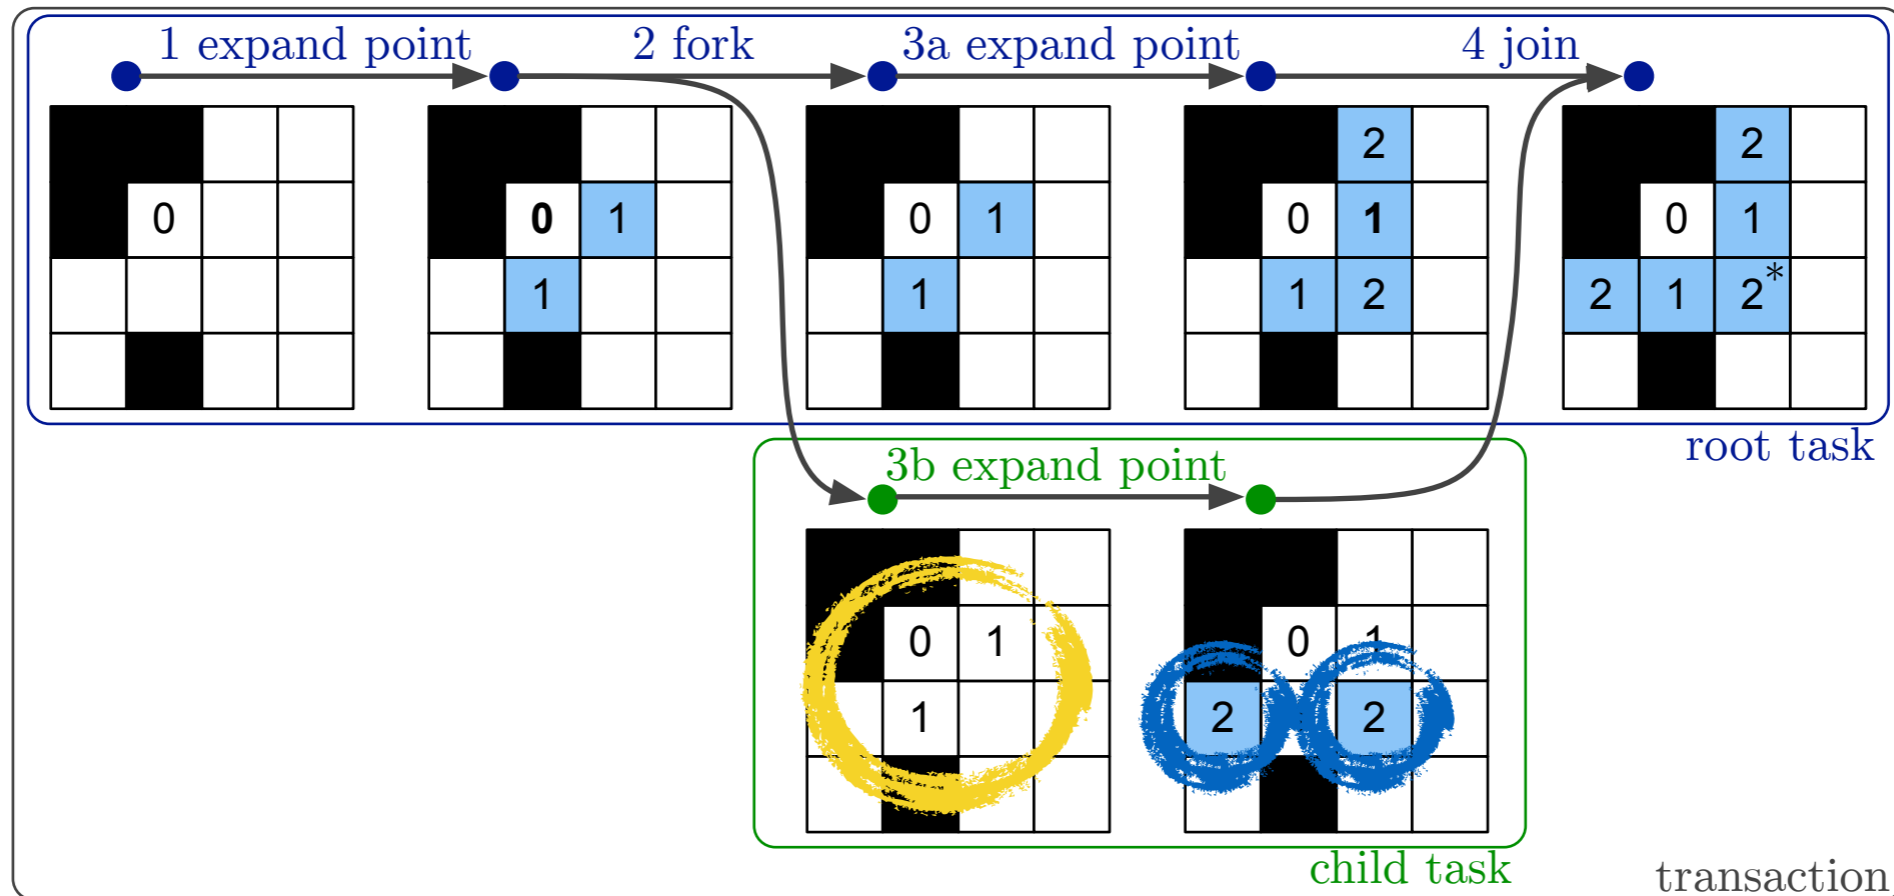
```
(atomic  
  (fork  
    (atomic ...)))
```

Transactional Tasks




- `fork`: copy the transactional state
- each task runs in isolation
- `join`: merge changes

fork



snapshot = tx state when created
local store = local modifications

join & conflicts



Merge
local
store

Conflict resolution function: (ref 0 resolve)

$\text{resolve} :: T \times T \times T \rightarrow T$

```
(defn resolve [o p c] c)
(defn resolve [o p c] p)
(defn resolve [o p c] (min p c))
(defn resolve [o p c] (+ p c))
(defn resolve [o p c] (error "Conflict on merge"))
```

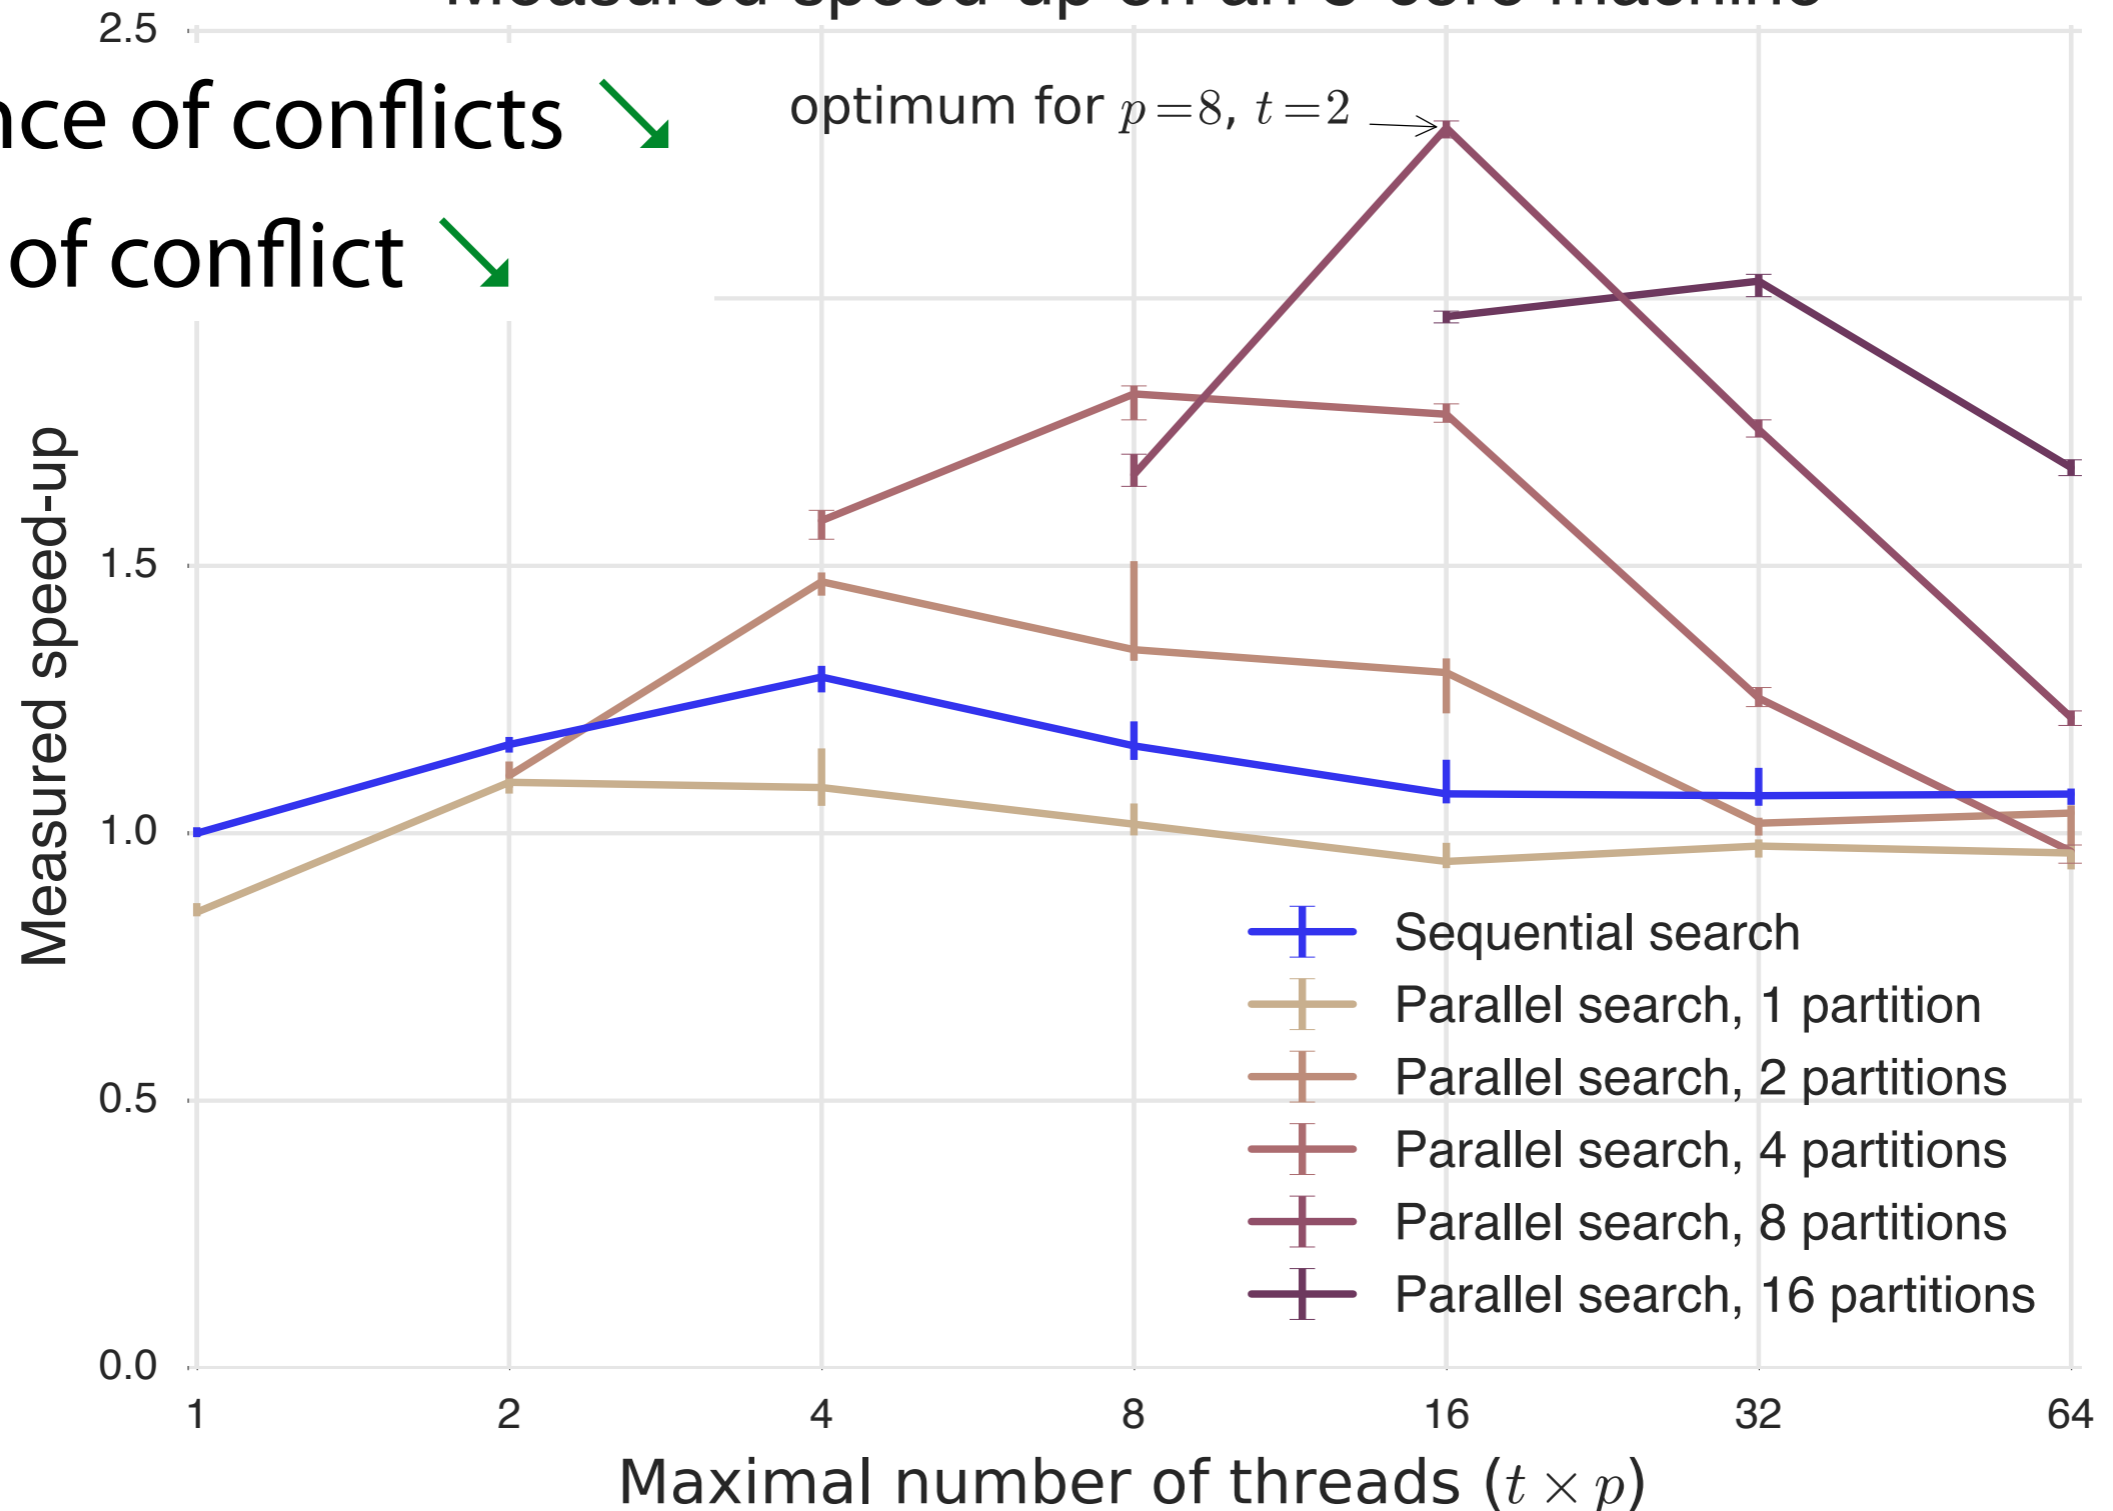
Properties

- **In-transaction parallelism** possible
- **Serializability** of transactions
- **Coordination** of tasks: all or none
- In-transaction **determinacy**
- Non-transactional tasks **unmodified**


Validation: Labyrinth

Measured speed-up on an 8-core machine

- Chance of conflicts ↘
- Cost of conflict ↘



Validation: Bayes



Take-away from experiments

- Labyrinth: parallelize search algorithm
⇒ fewer & cheaper conflicts
- Bayes: more fine-grained parallelism
- Low developer effort

Summary

Parallelism in a transaction is useful


But currently:

- ↗ not allowed (Haskell)
- ↘ not serializable (Clojure, Scala)

```
atomically $  
  do { forkIO ... }  
  
(atomic  
 (fork  
  (atomic ...)))
```

Idea: transactional tasks

- safe access to encapsulating tx
- 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/>

