#### **Concurrent Revisions**

Lecture 1: From Tasks to Revisions

#### Sebastian Burckhardt

**Microsoft Research** 

Joint work with Alexandro Baldassin, Daan Leijen (see also our papers in OOPSLA 2010, ESOP 2011)



#### Outline

- Parallel Programming With Tasks
  - Task Dags and Scheduling
  - Conflict Elimination
  - Motivation: Game Example
- Concurrent Revisions
  - Fork/Join and Revision Diagrams
- Semantics
  - Syntax and Semantics
  - Determinacy

# Task Programming Model

- Program expresses task dag (directed acyclic graph)
- Runtime dynamically schedules task dag onto multiprocessor
- Commonly used paradigm
   (CILK, Intel TBB, java.util.concurrent, Microsoft TPL, ...)
- Abstracts machine details (sync primitives, number of processors)
  - Often dynamic (graph generation and scheduling are interwoven)

#### Expressing parallelism

Example:

```
for (i = 0; i < n; i++)
{
    work1(i);
    work2(i);
}</pre>
```

#### Can parallelize loop iterations and/or loop body.

```
parallel for
  (i = 0; i < n; i++)
{
    work1(i);
    work2(i);
}</pre>
```

```
for
  (i = 0; i < n; i++)
  {
    parallel { work1(i), work2(i) }
  }
</pre>
```

What does it mean, exactly?

## **Comparison with TPL syntax**

```
Parallel.For(0, n, i =>
 {
    work1(i);
    work2(i);
 })
parallel for
(i = 0; i < n; i++)
{
   work1(i);
   work2(i);
}
```

```
for
 (i = 0; i < n; i++)
  Parallel.Invoke(
      () => { work1(i); },
     () => { work2(i); }
  });
}
for
 (i = 0; i < n; i++)
{
   parallel { work1(i), work2(i) }
}
```

# Task graphs



# What do we mean by task graphs?

- Vertices
  - Represent tasks or synchronization
- Directed Edges
  - Represent scheduling constraints
- No cycles
  - Thus, directed acyclic graph (dag)

Task graphs (a.k.a **task dags**) are an excellent visualization for understanding parallelism.



### The meaning of edges

- An edge v -> w means
   "task v must complete before task w starts"
- Is a transitive property!
   (a before b) and (b before c) implies
   (a before c)
- We omit edges that are implied by transitivity (take the "transitive reduction")

### Examples



for
 (i = 0; i < n; i++)
{
 parallel { work1(i), work2(i) }
}</pre>



#### **Executing Task Dags**

• On a single processor: serial schedules



Many schedules possible: ABCDEF, ABCEDF, ACBDEF, ACBEDF, ACDBEF, ACEBDF, ACDEBF, ACDEBF, ACEDBF, ..., CADEB, CAEDB, ... CDEABF, CEDABF

 All schedules have same execution time:

• 
$$T_1(G) = A + B + C + D + E + F$$

# Parallel Schedules



- Suppose we have an idealized parallel processor (infinite number of processors, no synchronization overhead)
- How long does it take to execute this task dag?

# **Parallel Schedules**



- Each task can start as soon as all previous tasks complete
- Execution time = critical path = time of longest path through the graph

### Work and Span

 We call the serial execution time T<sub>1</sub>(G) the work of G.

#### $T_1(G) = \sum \{ T(v) \mid v \text{ in } G \}$

 We call the ideal parallel execution time T<sub>∞</sub>(G) the span of G:

 $T_{\infty}(G) = \max \{ T_1(p) \mid p \text{ maximal path of } G \}$ 

# Speedup

• Divide work by span:

• What does the ratio  $T_{\infty}(G) / T_1(G)$  mean?

 It's how much we can improve the execution time by parallel execution! We call this par(G), parallelism of G.

It's an upper bound on the possible speedup.

# Example



- Work T<sub>1</sub>(G) is 4.
- Span  $T_{\infty}(G)$  is 3.
- Parallelism par(G) is 4/3
- Speed improvement from parallelism is never more than 33%
- Clear from picture: Using more than 2 processors has no benefits.

### Special case: Amdahl's law

- Don't remember formula? No problem, can derive it instantly.
  - p : parallel fraction of program
  - 1-p : sequential fraction of program
  - n : number of processors



- Work is 1
- Span is (1-p) + p/n
- Parallelism is work/span
   1 (1-p) + p/n

#### Execution time on n processors

 Define T<sub>n</sub>(G) to be the best possible execution time on n processors (that is, the smallest time we can achieve for any schedule of G on n processors)

• Clearly, we can give bounds  $T_{\infty}(G) \leq T_{n}(G) \leq T_{1}(G)$ 

#### FYI: How hard to find best schedule?

Finding the best schedule is hard

 Computing T<sub>n</sub>(G) is NP-complete for n>2 [Gary, Johnson 1979]

Finding a reasonably good schedule is easy

[Brent's Theorem] for any dag G,
 T<sub>n</sub>(G) ≤ T<sub>1</sub>(G)/n + T<sub>∞</sub>(G)

[Blumofe/Leiserson 1999]: Any greedy schedule achieves this bound.

#### Series-Parallel Task Dags

- We can build many task dags using just parallel and serial composition.
- These dags are called 'series-parallel'
- If our programs express parallelism using only the following constructs, the dags are always seriesparallel:
  - parallel { }
  - parallel for

# Serial Composition of Task Dags

Given two task dags G, H, build serial composition (G; H) by

 connecting all sinks of G to all sources

 To use fewer edges, we can construct (G; v; H) where v is an extra synchronization vertex (with T(v) = 0)





Parallel Composition of Task Dags

Given two task dags G, H, build parallel composition (G II H) by

Building disjoint union of G, H

Or, to keep dag connected, we can construct
 (s ; (G || H); e)
 where s,e are extra
 synchronization vertices
 (with T(s) = T(e) = 0)





### Examples



for
 (i = 0; i < n; i++)
{
 parallel { work1(i), work2(i) }
}</pre>



#### More ways to express parallelism

 Can use fork, join primitives to explicitly create parallel tasks

var t; t = fork { A(); B(); join t; C();

#### More ways to express parallelism

Can use fork, join primitives to explicitly create parallel tasks





#### **Comparison to TPL syntax**

Task t;
t = fork {
 A();
}
B();
join t;

```
Task t;
t = Task.Factory.StartNew(() =>
    {
        A();
    });
B();
t.Wait();
```

Our join is more restricted than wait... Must be called exactly once.

#### Can express parallel using fork, join

```
function
myparallelinvoke(f, g)
{
    var t = fork {
        f();
    }
    g();
    join t;
}
```

```
function myparallelfor(a, b, f)
{
    var t = new Task[n];
    for (int i = a; i < b; i++)
        t[i] = fork { f(i); }
    for (int i = a; i < b; i++)
        join t[i];
}</pre>
```

Can we do the opposite, i.e. express fork, join using parallel?

#### Not all dags are series-parallel

 For example, we can use asymmetric fork and join to create task graphs like this





#### Not all dags are series-parallel

 For example, we can use asymmetric fork and join to create task graphs like this





So, what does all this mean for how we should write parallel programs?

### Thus, for fastest execution...

• Make everything as parallel as possible.





 Unfortunately, very few programs are fully parallel ("embarassingly parallel" or "pleasantly parallel")
 Bocause tasks usually exhibit dependencies

Because tasks usually exhibit dependencies and conflicts!

#### Parallelism Blockers

#### Conflicts

- Two tasks A, B access the same data, and one (or both) modify it
- (If both tasks are only reading, we don't call it a conflict)
- Two tasks that don't have conflicts can run in parallel
- Dependence (a type of conflict)
  - A conflict as above, where A writes some result and B reads that result

#### Terminology: Conflict vs. Data Race

- All data races are conflicts.
- Not all conflicts are data races.



- Has conflict.
- Has data race.

- Has conflict.
- does NOT have data race.

#### Hazards

• Sequential execution is **not equivalent** to parallel execution if there are *hazards*.

RAW hazard (read after write) WAR hazard (write after read)

WAW hazard (write after write)

a.k.a. data dependency a.k.a. Anti-dependency

#### Hazards

- Hazards in sequential composition correspond to *conflicts* in parallel composition.
- Conflict = concurrent tasks access shared data, and at least one of them modifies it.



Read-Write conflict

Write-write conflict

Read-Write conflict

### Why are conflicts bad?

- Nondeterminism
  - The order of the accesses affects the outcome
- Data Consistency
  - Many typical data representations (e.g. linked list) are not safe unless carefully crafted (see M.H.) or protected by locks

#### Performance

 Cache coherence: Conflicts cause cache misses (shared memory is not actually shared)

#### Why are conflicts bad?

- Nondeterminism
- Data Consistency
- Performance

Essential insight:

Eliminate frequent conflicts. Make rare conflicts safe.
### Some conflicts are superfluous.

• WAR and WAW can usually be eliminated.



- Because they are not true dependencies but a matter of reading/writing correct versions
- work2 can proceed without work1's result if
  - It does not interfere with work1's input (RAW)
  - It ensures to overwrite work1's output (WAW)

### Some conflicts are superfluous.

 For example, we can eliminate the following WAR / WAW hazards (privatization).



- Trick: copying / renaming
  - Register renaming extremely effective in hardware: Tomasulo's algorithm

### The Art of Task-Parallel Programming

- Identify tasks to run in parallel
  - Make tasks large enough
  - Make tasks small enough
  - Make tasks parallel enough
- Carefully consider conflicts
  - Eliminate or reduce conflicts
    - Tricks of the trade
  - Make remaining conflicts safe
    - Concurrency control

### The Art of Task-Parallel Programming

- Identify tasks to run in parallel
  - Make tasks large enough
  - Make tasks small enough
  - Make tasks parallel enough
- Carefully consider conflicts
  - Eliminate or reduce conflicts
    - Tricks of the trade
  - Make remaining conflicts safe
    - Concurrency control

QUANTITATIVE CRITERIA depend on characteristics of workload and machine

QUALITATIVE CRITERIA Do not depend on machine or workload

## Make tasks small enough

total execution time >= time for longest task

# Make tasks parallel enough

total execution time >= time along critical path

(= path from source to sink whose sum of task times is maximal)

## Make tasks large enough

### • Danger:

- Processors are very fast at sequential execution.
- Processors are slow at scheduling/starting/ending tasks.
- Unless a task contains a significant amount of work, it is not worth to try parallel execution!
- Usually not a problem when using parallel for and lots of iterations
  - Smart implementation of parallel for chunks iterations automatically and dynamically

## **Example: Histogram Computation**

- Suppose f is a function with range 0..7
- We want to count how many times each number in 0..7 is produced when calling f on 0...n-1
- We assume n large enough to make parallelism worthwhile (e.g. n = 1000000)

## Simple but incorrect solution

 Does not count correctly because ++ is not atomic

### **Correct but slow solution**

```
a = new int[8];
l = new lock[8];
parallel for
  (i = 0; i < n; i++)
{
     var r = f(i);
     lock( l[r] )
        { a[r]++; }
}
```

```
a = new int[8];
parallel for
   (i = 0; i < n; i++)
{
    var r = f(i);
    atomic-increment(&(a[r]));
}</pre>
```

- Atomic increment is typically a bit faster than lock/unlock.
- Both solutions are correct, but overall performance is quite bad (worse than sequential)

### A better solution



- We have eliminated the conflicts.
- Worse in theory (task dag), but much better in practice

### The toolbox

Eliminate frequent conflicts

- Architectural Patterns
  - Producer-Consumer
  - Pipeline
  - Worklist
  - ....
- Replication Patterns
  - Immutability
  - Double Buffering
  - Concurrent Revisions

Make rare conflicts safe

- pessimistic concurrency control
   coarse- or fine-grained locking
- optimistic concurrency control

speculate on absence of conflicts, roll back if speculation fails

Transactional memory

### PART II: INTRODUCTION TO CONCURRENT REVISIONS

Targeted Application Scenario Shared Data and Parallel Tasks

Reader	Mutat	or	
		Reader	
Shared Data			
	Reader	]	
Mutator	Autator Mutator		

- Shared state is read/mutated by many tasks.
- Both inter- and intra-task parallelism.
- Challenge: <u>Tasks exhibit</u> frequent conflicts.

## 3 Examples of this Pattern: Office Browser



Game





## SpaceWars Game



#### Sequential Game Loop:

#### while (!done)

input.GetInput(); input.ProcessInput(); physics.UpdateWorld(); for (int i = 0; i (physics.numsplits; i++) physics.CollisionCheck(i); network.SendNetworkUpdates(); network.HandleQueuedPackets(); if (frame % 100 = 0) SaveGame(); ProcessGuiEvents(); screen.RenderFrameToScreen(); audio.PlaySounds(); frame++;

## Can you parallelize this loop?

while (!done)	Conflicts on object
{	Coordinator
input.GetInput();	Coordinates.
input.ProcessInput();	
physics.UpdateWorld();	Reads and writes all positions
for (int i = 0; i (physics.numsplits; i++)	
physics.CollisionCheck(i);	Writes some positions
network.SendNetworkUpdates();	
network.HandleQueuedPackets();	→ Writes some positions
if (frame % 100 = 0)	
SaveGame():	
ProcessGuiEvents()	redus all positions
screen RenderFrameToScreen()	. Deede ell restitions
audio PlaySounde():	Reads all positions
audio.FlaySourius(),	
Trame++;	
}	

## Fundamental Insights

- Intra-loop RW conflicts can be eliminated: Tasks do not need most recent version of state. What they need is a snapshot of the state, taken at the beginning of each iteration.
- Intra-loop WW conflicts can be eliminated: All we need to do is get the priority right (network task trumps conflict detection task trumps physics task)
- How to express this intent? Not equivalent to *any* sequential execution of the tasks.

### **Concurrent Revisions Model**



- fork and join revisions (lines with arrow at end)
- Revisions are isolated
  - fork copies all state
  - join replays updates along the arrow at the tip of the arrow
- Use named operations (add/set) instead of basic assignment

## **Revision Diagrams vs. Task Dags**

- Different look
  - Rounded corners
  - No boxes for fork, join
  - Arrow tip only at end of revision
- Less general
  - ONLY asymmetric fork, join
  - no symmetric parallel

## Side-By-Side Example





### **Understanding Concurrent Revisions**

• Fork copies the current state.

 Join replays updates at the tip of arrow.



 State determined by sequence of updates along path from root



- State determined by sequence of updates along path from root
- Join replays updates at the tip of arrow.







## Applying this idea

 So now, can we express the intended parallelization of the SpaceWars Loop?

- Intra-loop RW conflicts can be eliminated: Tasks do not need most recent version of state. What they need is a snapshot of the state, taken at the beginning of each iteration.
- Intra-loop WW conflicts can be eliminated: All we need to do is get the priority right (network task trumps conflict detection task trumps physics task)

### **Revision Diagram of Parallelized Game Loop**



### **Eliminated Read-Write Conflicts**



### **Eliminated Write-Write Conflicts**





- Autosave now perfectly unnoticeable in background
- Overall Speed-Up: 3.03x on four-core (almost completely limited by graphics card)



### PART III: A LOOK AT CODE

## The C# library & IL rewriter

- We have seen the programming model
  - versioned data
  - fork/join of revisions
- But what does it concretely look like?
  Let's see some code.

Or do it yourself: visit <u>http://rise4fun.com</u> and play

### **Version Data**

```
[Versioned] int z;
[Versioned] string s = "abc";
```

}

```
[Versioned, MergeWith("AdditiveMerge")]
int i = 0;
```

## Using Simple Fork/Join


# **Ordering Joins**

RevisionTask a = CurrentRevision.Fork(() => x = 1;}); RevisionTask b = CurrentRevision.Fork(() => ł x = 2;}); CurrentRevision.Join(b); CurrentRevision.Join(a); Console.WriteLine("x = " + x ); // writes 1

# **Abandoning Revisions**

- We can abandon a revision (instead of join)
  - Discards the version once task completes
  - Is not a cancellation: task still runs to completion



## Example: Background Save

- If a task does only read from shared state, we can abandon it to avoid blocking on a join
- Example: background save operation
  - never blocks application
  - snapshot survives as long as it needs to

RevisionTask backgroundsave =
 CurrentRevision.Fork(() => { SaveStateToFile(); });

// don't wait for task... let it finish on its own time
CurrentRevision.Abandon(backgroundsave);

# PART IV: FORMALIZATION & DETERMINACY

**ADVANCED SECTION – BONUS MATERIAL** 

## **Formal Semantics**

- See paper in [ESOP 2011]
- Similar to the AME calculus by Abadi et al.
- Proof of determinism
- Formal correspondence to the Revision Diagrams
- Proof of the semi-lattice property
- Can be shown to generalize 'snapshot isolation'

## Syntax

#### Syntactic Symbols $v \in Val$ $::= c \mid x \mid l \mid r \mid \lambda x.e$ $c \in Const ::= unit | false | true$ $l \in Loc$ $r \in Rid$ $x \in Var$ $e \in Expr$ ::= v $e \ e \ | \ (e \ ? \ e \ : e)$ ref e | !e | e := e rfork e | rjoin e

#### State

*LocalStore*  $= Loc \rightarrow Val$ 



#### State

 $s \in GlobalState = Rid \rightarrow LocalState$ *LocalState* = *Snapshot* × *LocalStore* × *Expr*  $\sigma \in Snapshot = Loc \rightarrow Val$  $\tau \in LocalStore = Loc \rightarrow Val$ **m**2 **m**3 { m->({}, {}, m0) } { m->({}, {x->0}, m1) } m2  $\{m > (\{\}, \{x > 0\}, m2), a > (\{x > 0\}, \{\}, a0) \}$  $\{m > (\{\}, \{x > 0\}, m2), a > (\{x > 0\}, \{x > 3\}, a1) \}$  $\{m > (\{\}, \{x > 0\}, m2), a > (\{x > 0\}, \{x > 3\}, a2), b > (\{x > 3\}, \{\}, b0) \}$  $\{m > (\{\}, \{x > 0\}, m2), a > (\{x > 0\}, \{x > 4\}, a3), b > (\{x > 3\}, \{\}, b0) \}$  $\{m > (\{\}, \{x > 0\}, m2), a > (\{x > 0\}, \{x > 4\}, a3), b > (\{x - >3\}, \{x - >5\}, b1) \}$  $\{ m > (\{\}, \{x > 4\}, m3), \}$ 

 $\{ m > (\{\}, \{x > 6\}, m4) \}$ 



#### **Execution Contexts**

$$\begin{split} \mathcal{E} &= \Box \\ &\mid \mathcal{E} \; e \mid v \; \mathcal{E} \mid (\mathcal{E} \, ? \, e : e) \\ &\mid \mathsf{ref} \; \mathcal{E} \mid ! \mathcal{E} \mid \mathcal{E} := e \mid l := \mathcal{E} \\ &\mid \mathsf{rjoin} \; \mathcal{E} \end{split}$$

 An expression matches at most one redex

#### Redexes

$$\begin{array}{l} \mathcal{E}[(\lambda x.e) \; v] \\ \mathcal{E}[(\mathsf{true} ? \; e_1 : e_2)] \\ \mathcal{E}[(\mathsf{false} ? \; e_1 : e_2) \\ \mathcal{E}[\mathsf{ref} \; v] \\ \mathcal{E}[!l] \\ \mathcal{E}[l := v] \\ \end{array}$$

#### Expressions

$$v \in Val \quad ::= c \mid x \mid l \mid r \mid \lambda x.e$$
  

$$c \in Const ::= unit \mid false \mid true$$
  

$$l \in Loc$$
  

$$r \in Rid$$
  

$$x \in Var$$
  

$$e \in Expr \quad ::= v$$
  

$$\mid e \mid e \mid (e ? e : e)$$
  

$$\mid ref \mid e \mid e \mid e := e$$
  

$$\mid rfork \mid e \mid rjoin \mid e$$

## **Operational Semantics**

$$\begin{split} s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\lambda x.e) \ v] \rangle) \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\mathsf{true} \ ? \ e_1 : e_2)] \rangle) \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\mathsf{false} \ ? \ e_1 : e_2)] \rangle) \end{split}$$

 $\begin{aligned} s(r \mapsto \langle \sigma, \tau, \mathcal{E}[\mathsf{ref } v] \rangle) \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[!l] \rangle) \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[l \coloneqq v] \rangle) \end{aligned}$ 

 $\begin{array}{l} \rightarrow_r s[r \mapsto \langle \sigma, \tau, \mathcal{E}[[v/x]e] \rangle] \\ \rightarrow_r s[r \mapsto \langle \sigma, \tau, \mathcal{E}[e_1] \rangle] \\ \rightarrow_r s[r \mapsto \langle \sigma, \tau, \mathcal{E}[e_2] \rangle] \end{array}$ 

 $\begin{array}{l} \rightarrow_r s[r \mapsto \langle \sigma, \tau[l \mapsto v], \mathcal{E}[l] \rangle] \\ \rightarrow_r s[r \mapsto \langle \sigma, \tau, \mathcal{E}[(\sigma::\tau)(l)] \rangle] \\ \rightarrow_r s[r \mapsto \langle \sigma, \tau[l \mapsto v], \mathcal{E}[\mathsf{unit}] \rangle] \end{array}$ 

 $\begin{array}{ll} s(r \mapsto \langle \sigma, \tau, \mathcal{E}[\mathsf{rfork} \ e] \rangle) & \longrightarrow_r s[r \mapsto \langle \sigma, \tau, \mathcal{E}[r'] \rangle][r' \mapsto \langle \sigma :: \tau, \epsilon, e \rangle] \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[\mathsf{rjoin} \ r'] \rangle)(r' \mapsto \langle \sigma', \tau', v \rangle) \to_r s[r \mapsto \langle \sigma, \tau :: \tau', \mathcal{E}[\mathsf{unit}] \rangle][r' \mapsto \bot] \\ s(r \mapsto \langle \sigma, \tau, \mathcal{E}[\mathsf{rjoin} \ r'] \rangle)(r' \mapsto \bot) & \longrightarrow_r \epsilon \end{array}$ 

 $\begin{array}{ll} s \in GlobalState = Rid \rightharpoonup LocalState \\ LocalState = Snapshot \times LocalStore \times Expr \\ \sigma \in Snapshot = Loc \rightharpoonup Val \\ \tau \in LocalStore = Loc \rightharpoonup Val \end{array}$ 

## Local Operations w/o side effects

$$s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\lambda x.e) \ v] \rangle)$$
  

$$s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\mathsf{true} ? \ e_1 : e_2)] \rangle)$$
  

$$s(r \mapsto \langle \sigma, \tau, \mathcal{E}[(\mathsf{false} ? \ e_1 : e_2)] \rangle)$$

For some revision r, with snapshot  $\sigma$  and local modifications  $\tau$ 

 $s \in GlobalState = Rid \rightarrow LocalState$   $LocalState = Snapshot \times LocalStore \times Expr$   $\sigma \in Snapshot = Loc \rightarrow Val$   $\tau \in LocalStore = Loc \rightarrow Val$ 



 $\tau \in \mathit{LocalStore} = \mathit{Loc} \rightharpoonup \mathit{Val}$ 



 $\begin{array}{ll} s \in GlobalState = Rid \rightharpoonup LocalState \\ LocalState = Snapshot \times LocalStore \times Expr \\ \sigma \in Snapshot = Loc \rightharpoonup Val \\ \tau \in LocalStore = Loc \rightharpoonup Val \end{array}$ 

## Determinacy

 Execution is determinate! That is, the final state of program execution does not depend on any scheduling decisions.

**Theorem 1** (Determinacy). Let e be a program expression, and let  $e \downarrow s$  and  $e \downarrow s'$ . Then  $s \approx s'$ .

- With two caveats:
  - States could differ in identifiers used for revisions and locations. We say s ≈ s' if states s, s' are equivalent in this sense.
  - It is possible for some executions to diverge while others don't.

# Proving Determinacy (1/3)

 First: if we fix a particular revision r to take a step, the outcome is determinate

Lemma 1 (Local Determinism). If  $s_1 \approx s'_1$  and  $s_1 \rightarrow_r s_2$  and  $s'_1 \rightarrow_r s'_2$ , then  $s_2 \approx s'_2$ .

Because the expression representing the program can evaluate in at most one way

- An expression matches at most one redex
- Each redex matches a unique operational rule

# Proving Determinacy (2/3)

 Local Confluence: if two different revisions take a step, they both can take (zero or one) additional step to get to an equivalent state

**Lemma 2** (Strong Local Confluence). Let  $s_1$  and  $s'_1$  be reachable states that satisfy  $s_1 \approx s'_1$ . Then, if  $s_1 \rightarrow_r s_2$  and  $s'_1 \rightarrow_{r'} s'_2$ , then there exist equivalent states  $s_3 \approx s'_3$  such that both  $s_2 \rightarrow_{r'}^? s_3$  and  $s'_2 \rightarrow_r^? s'_3$ .

Reason: steps by different revisions do commute because they either

- Do not influence each other
- Are mutually exclusive (two revisions trying to join the same revision) thus triggering a transition to the error state.