

Correctness of Dynamic Dependence Analysis for Implicitly Parallel Tasking Systems

Wonchan Lee
Stanford University

George Stelle
LANL

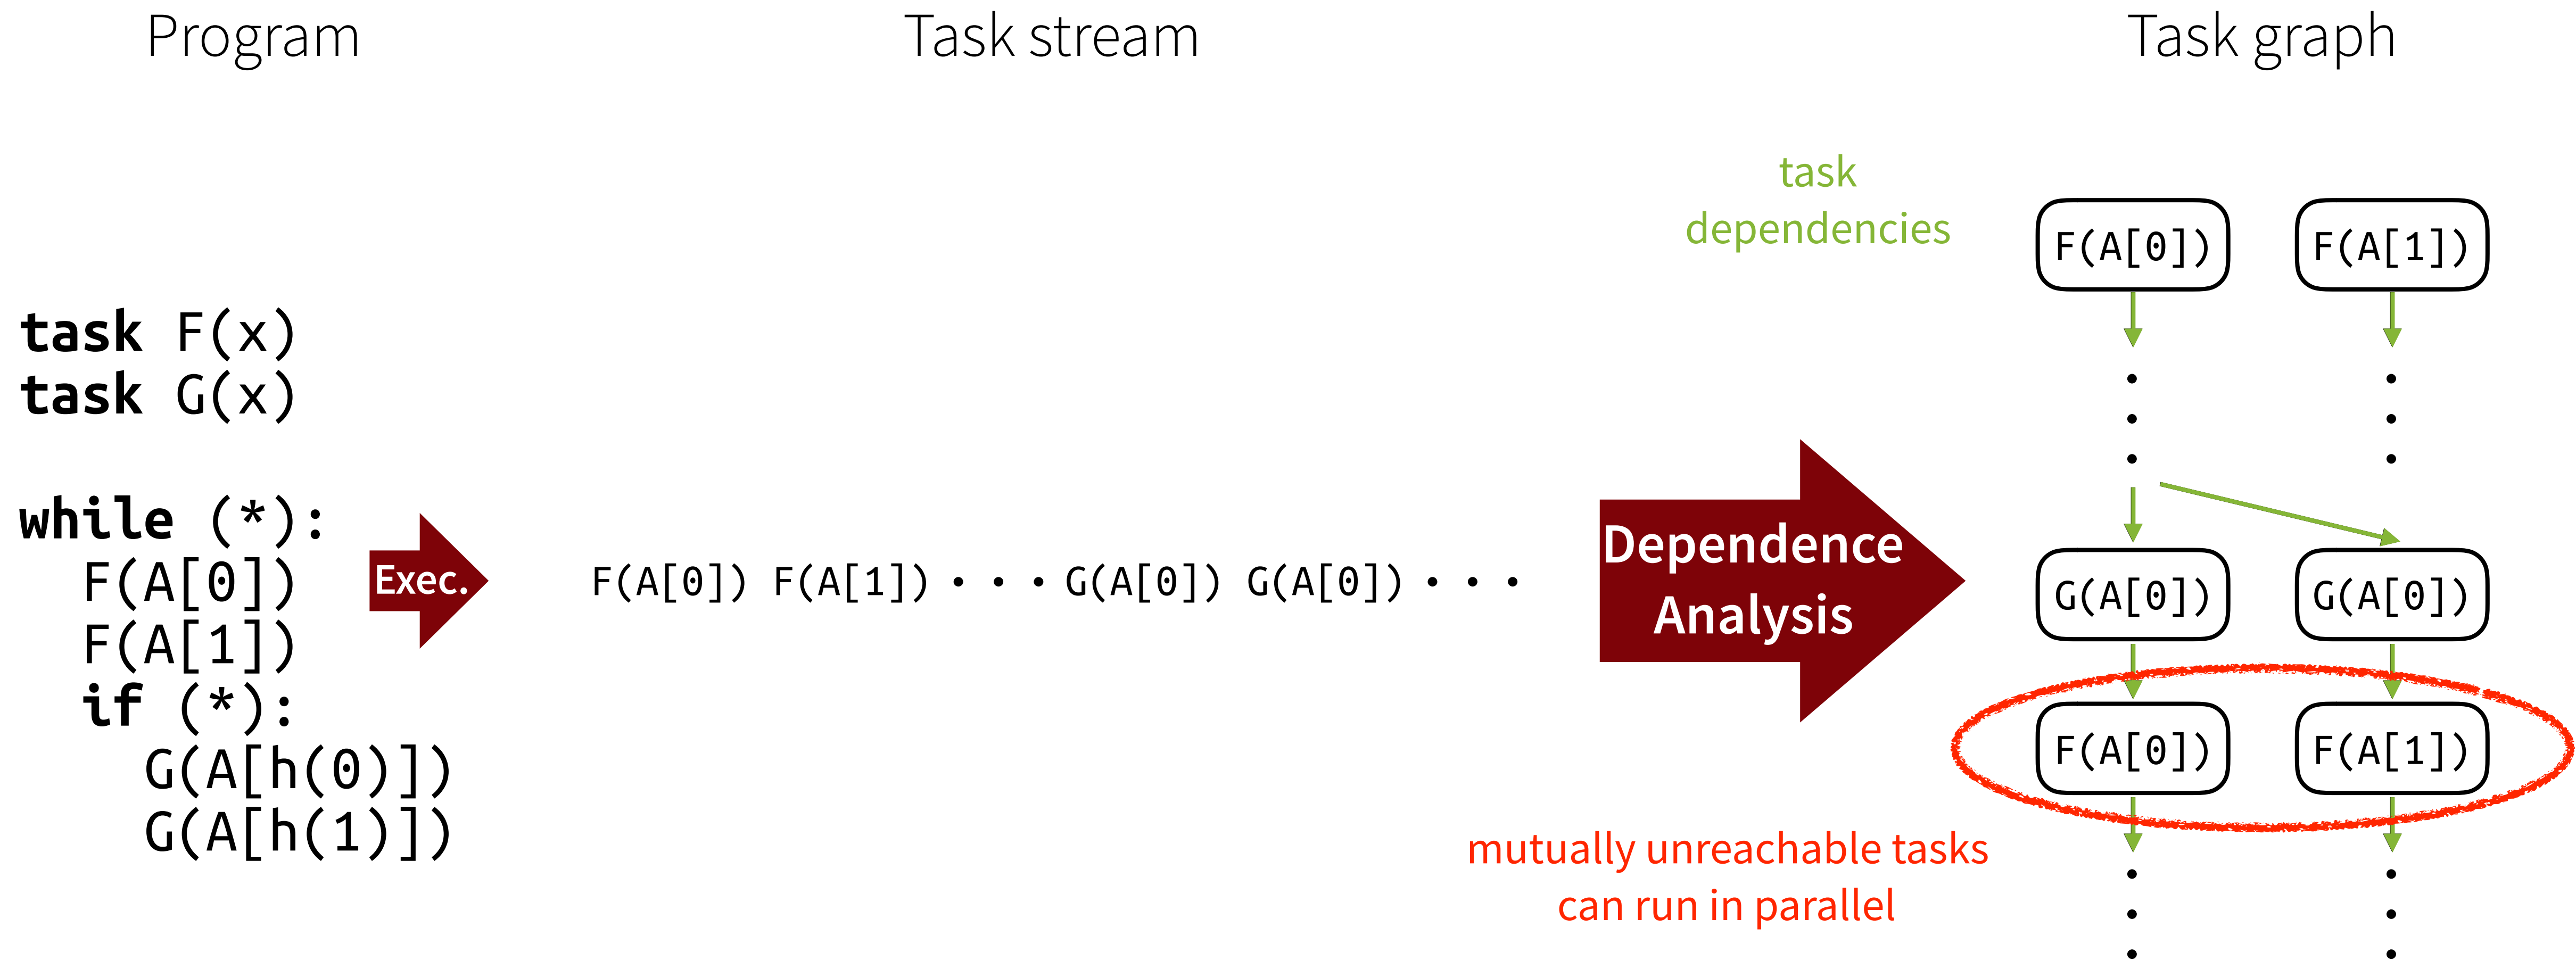
Patrick McCormick
LANL

Alex Aiken
Stanford University

November 12, 2018 @ Correctness'18

Implicitly Parallel Tasking Systems

- Automatically extract parallelism from program using dynamic dependence analysis



- Legion, PaRSEC, StarPU, PyTorch, etc.

Benefits of Implicitly Parallel Tasking Systems


- Sequential semantics
 - No parallel programming errors
- Precise task dependencies in the presence of dynamic behaviors

control divergence

```
task F(x)
task G(x)

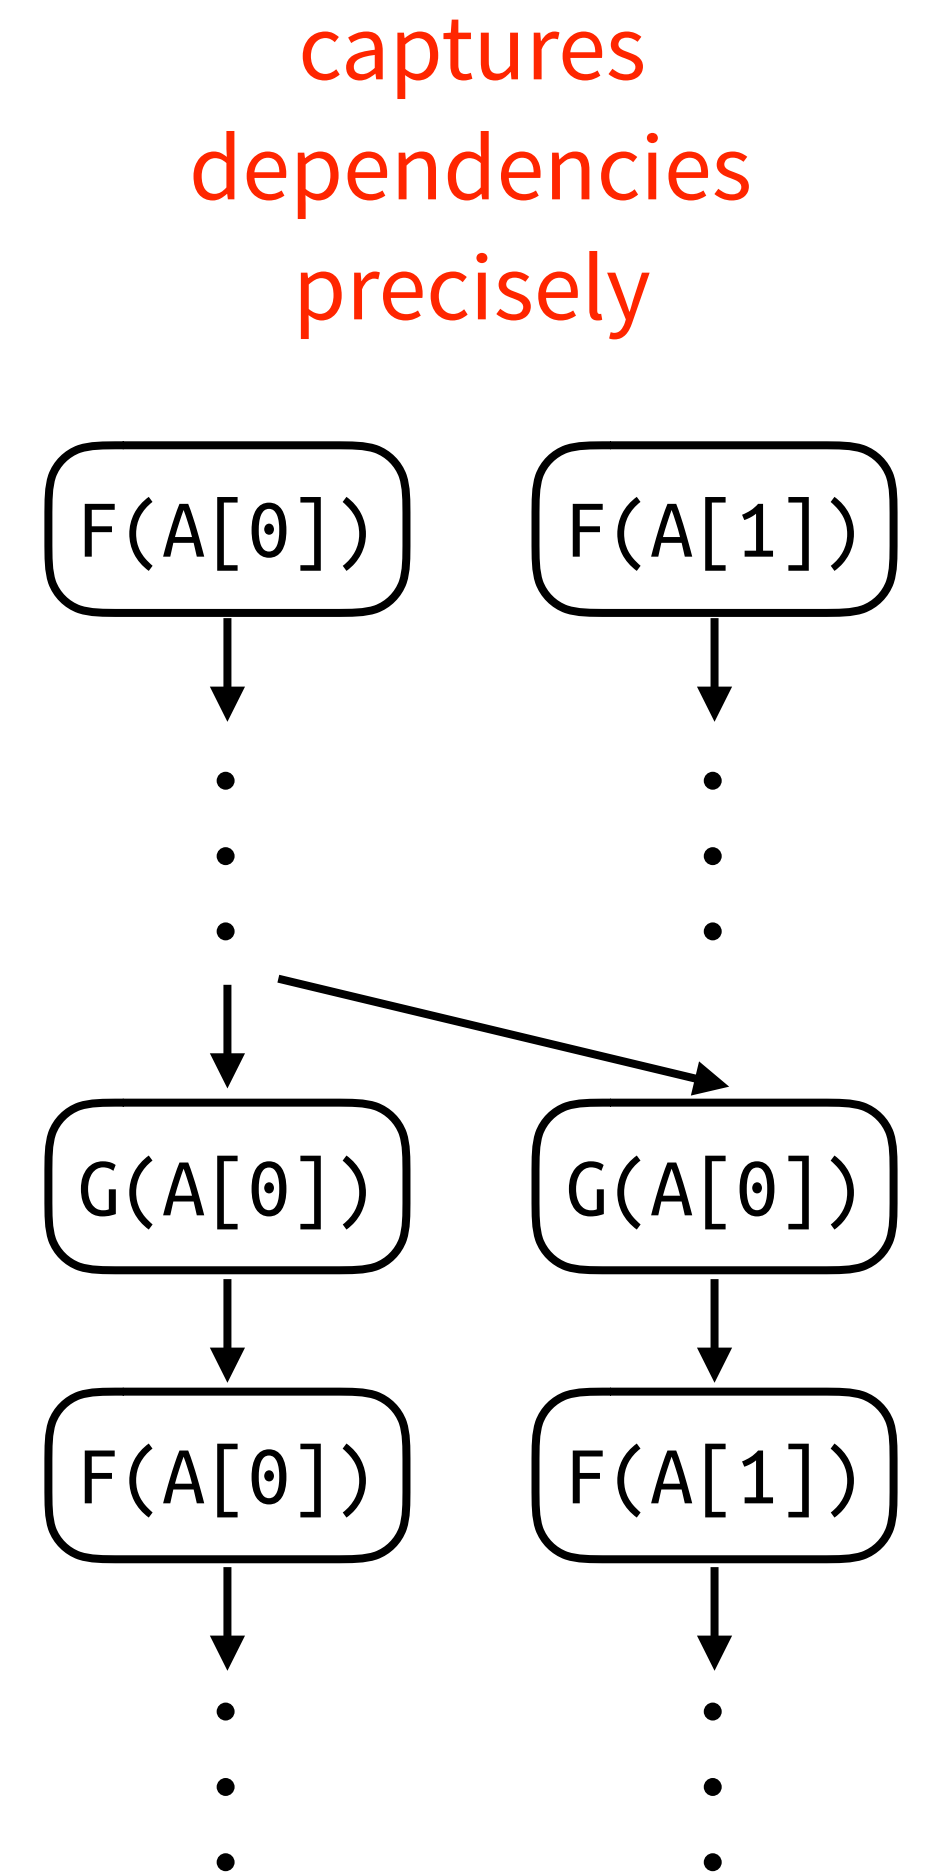

while (*):
  F(A[0])
  F(A[1])
  if (*):
    G(A[h(0)])
    G(A[h(1)])
```

opaque function calls



F(A[0]) F(A[1]) . . . G(A[0]) G(A[0]) . . .

no conditionals or function calls



Correctness of Dynamic Dependence Analysis

- Important
- Non-trivial to achieve
 - Due to optimizations in real algorithms
- Topic of this paper
 - D_{epoch} : a model dependence analysis algorithm using *epochs*
 - Correctness proof for D_{epoch}

Soundness and Completeness



- Soundness

- All dependencies are discovered

- Completeness


- No independent tasks are serialized unnecessarily

Task stream: WR(A) WR(B) RD(A) RD(B)

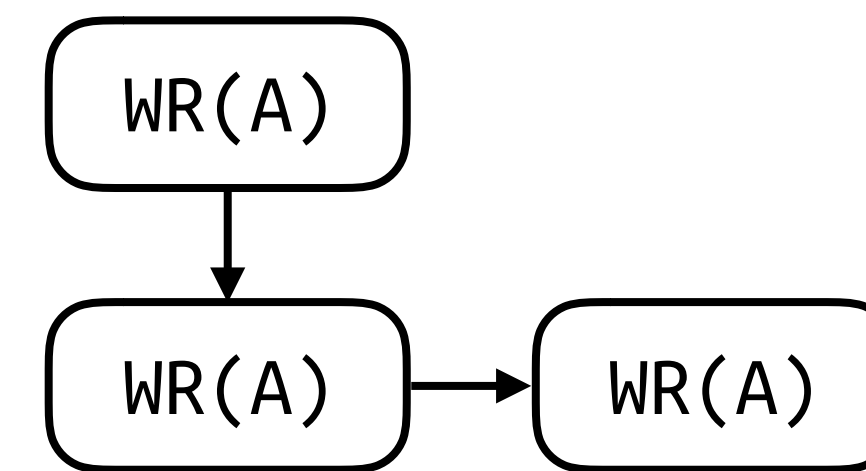


Bonus Point: Parsimony

- Parsimonious task graph
 - One with fewest edges among equivalent task graphs



has **transitive edge**
→ not parsimonious 😞



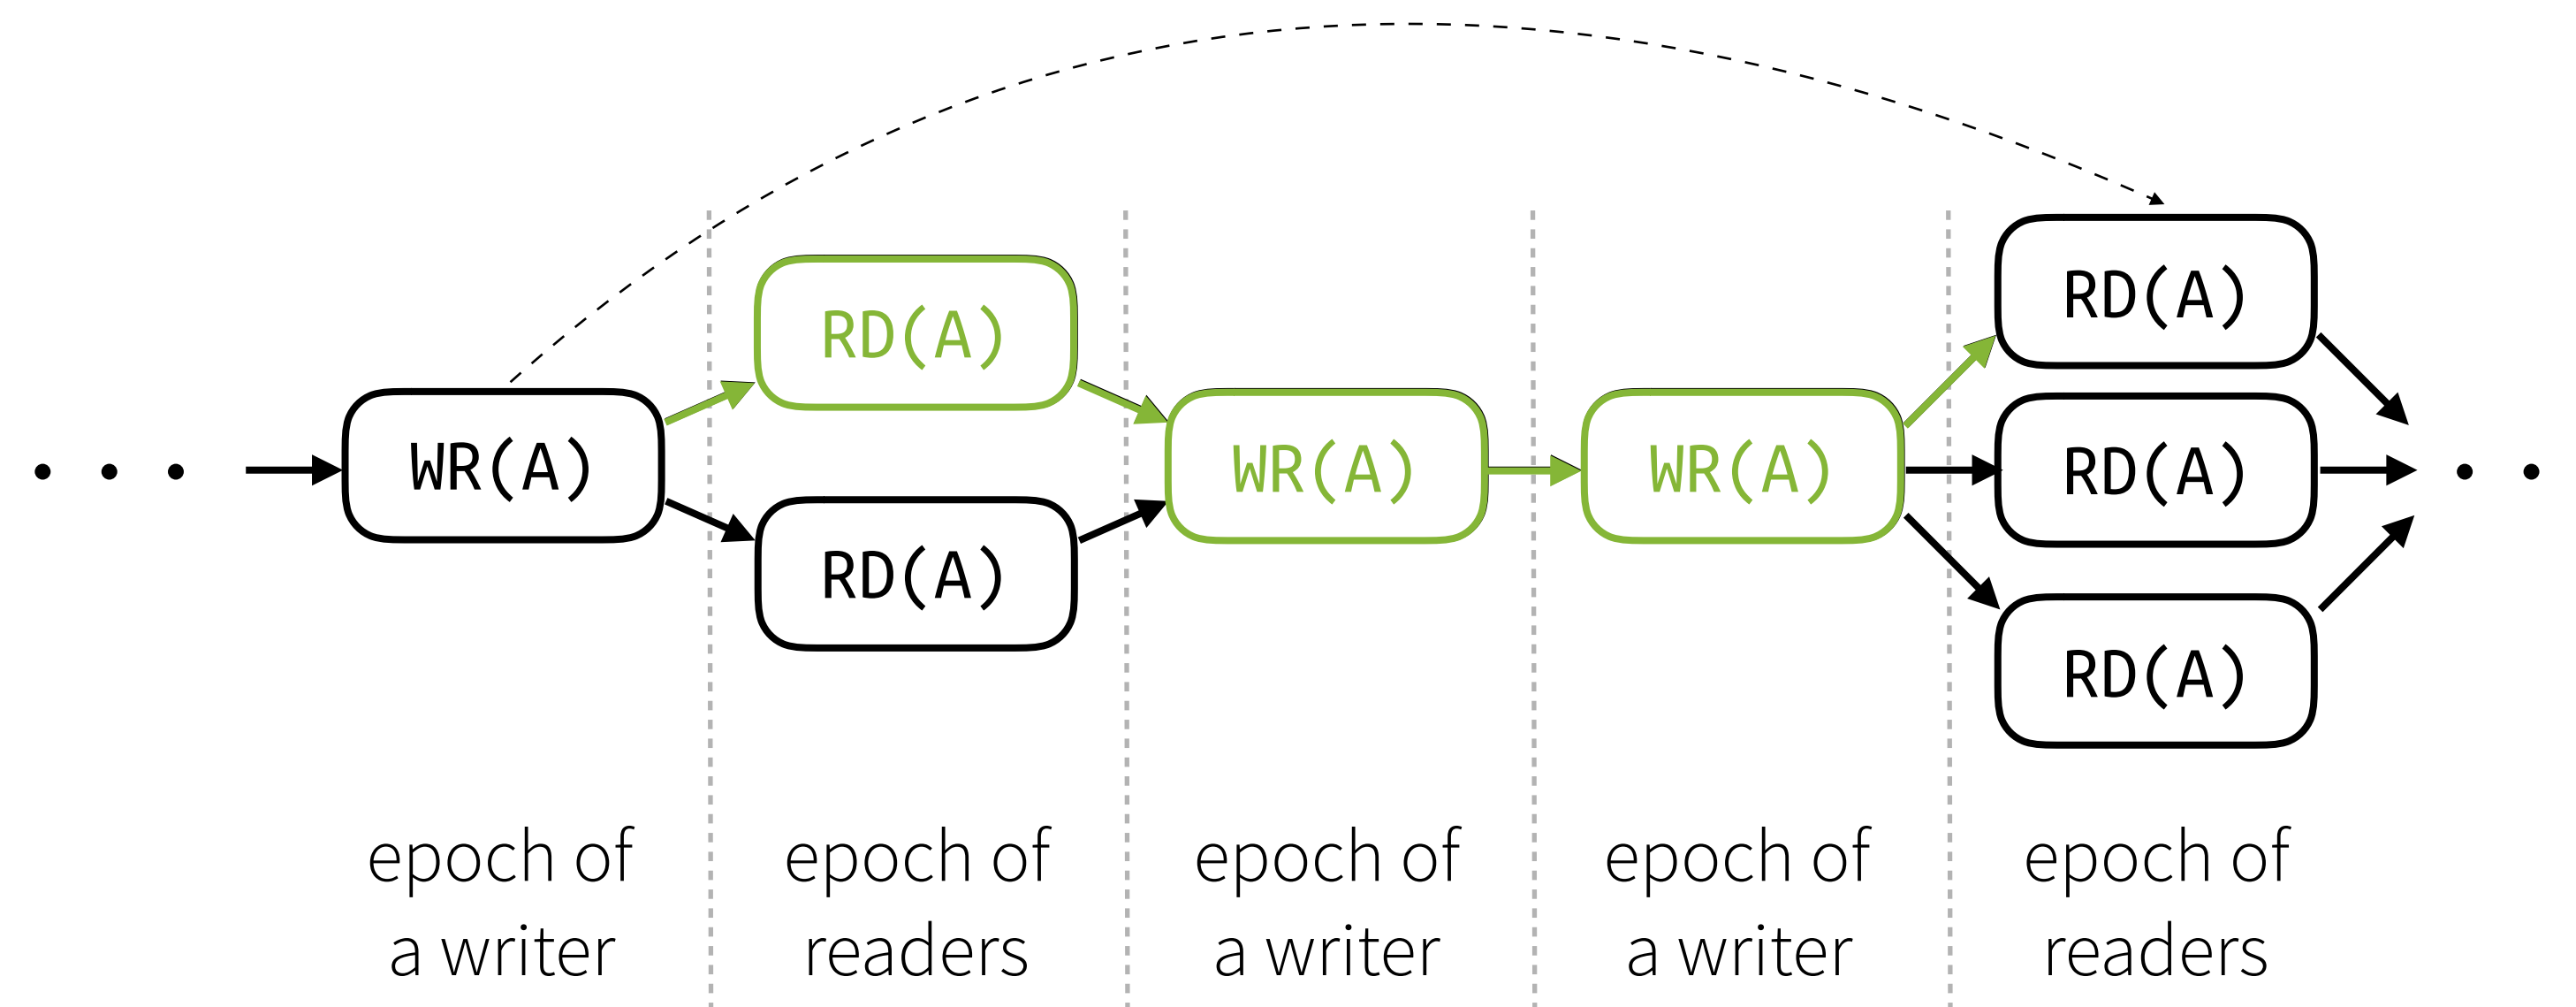
no transitive edge
→ **parsimonious** 👍

- Task graphs are often not parsimonious, because they require costly transitive reduction
 - Still valuable to understand when algorithms produce parsimonious task graphs

D_{epoch} : Epoch-based Algorithm

- Based on SWMR invariant: either **Single Writer** or **Multiple Readers** can exist for a region
- Tasks accessing a single region form *epochs* of independent tasks

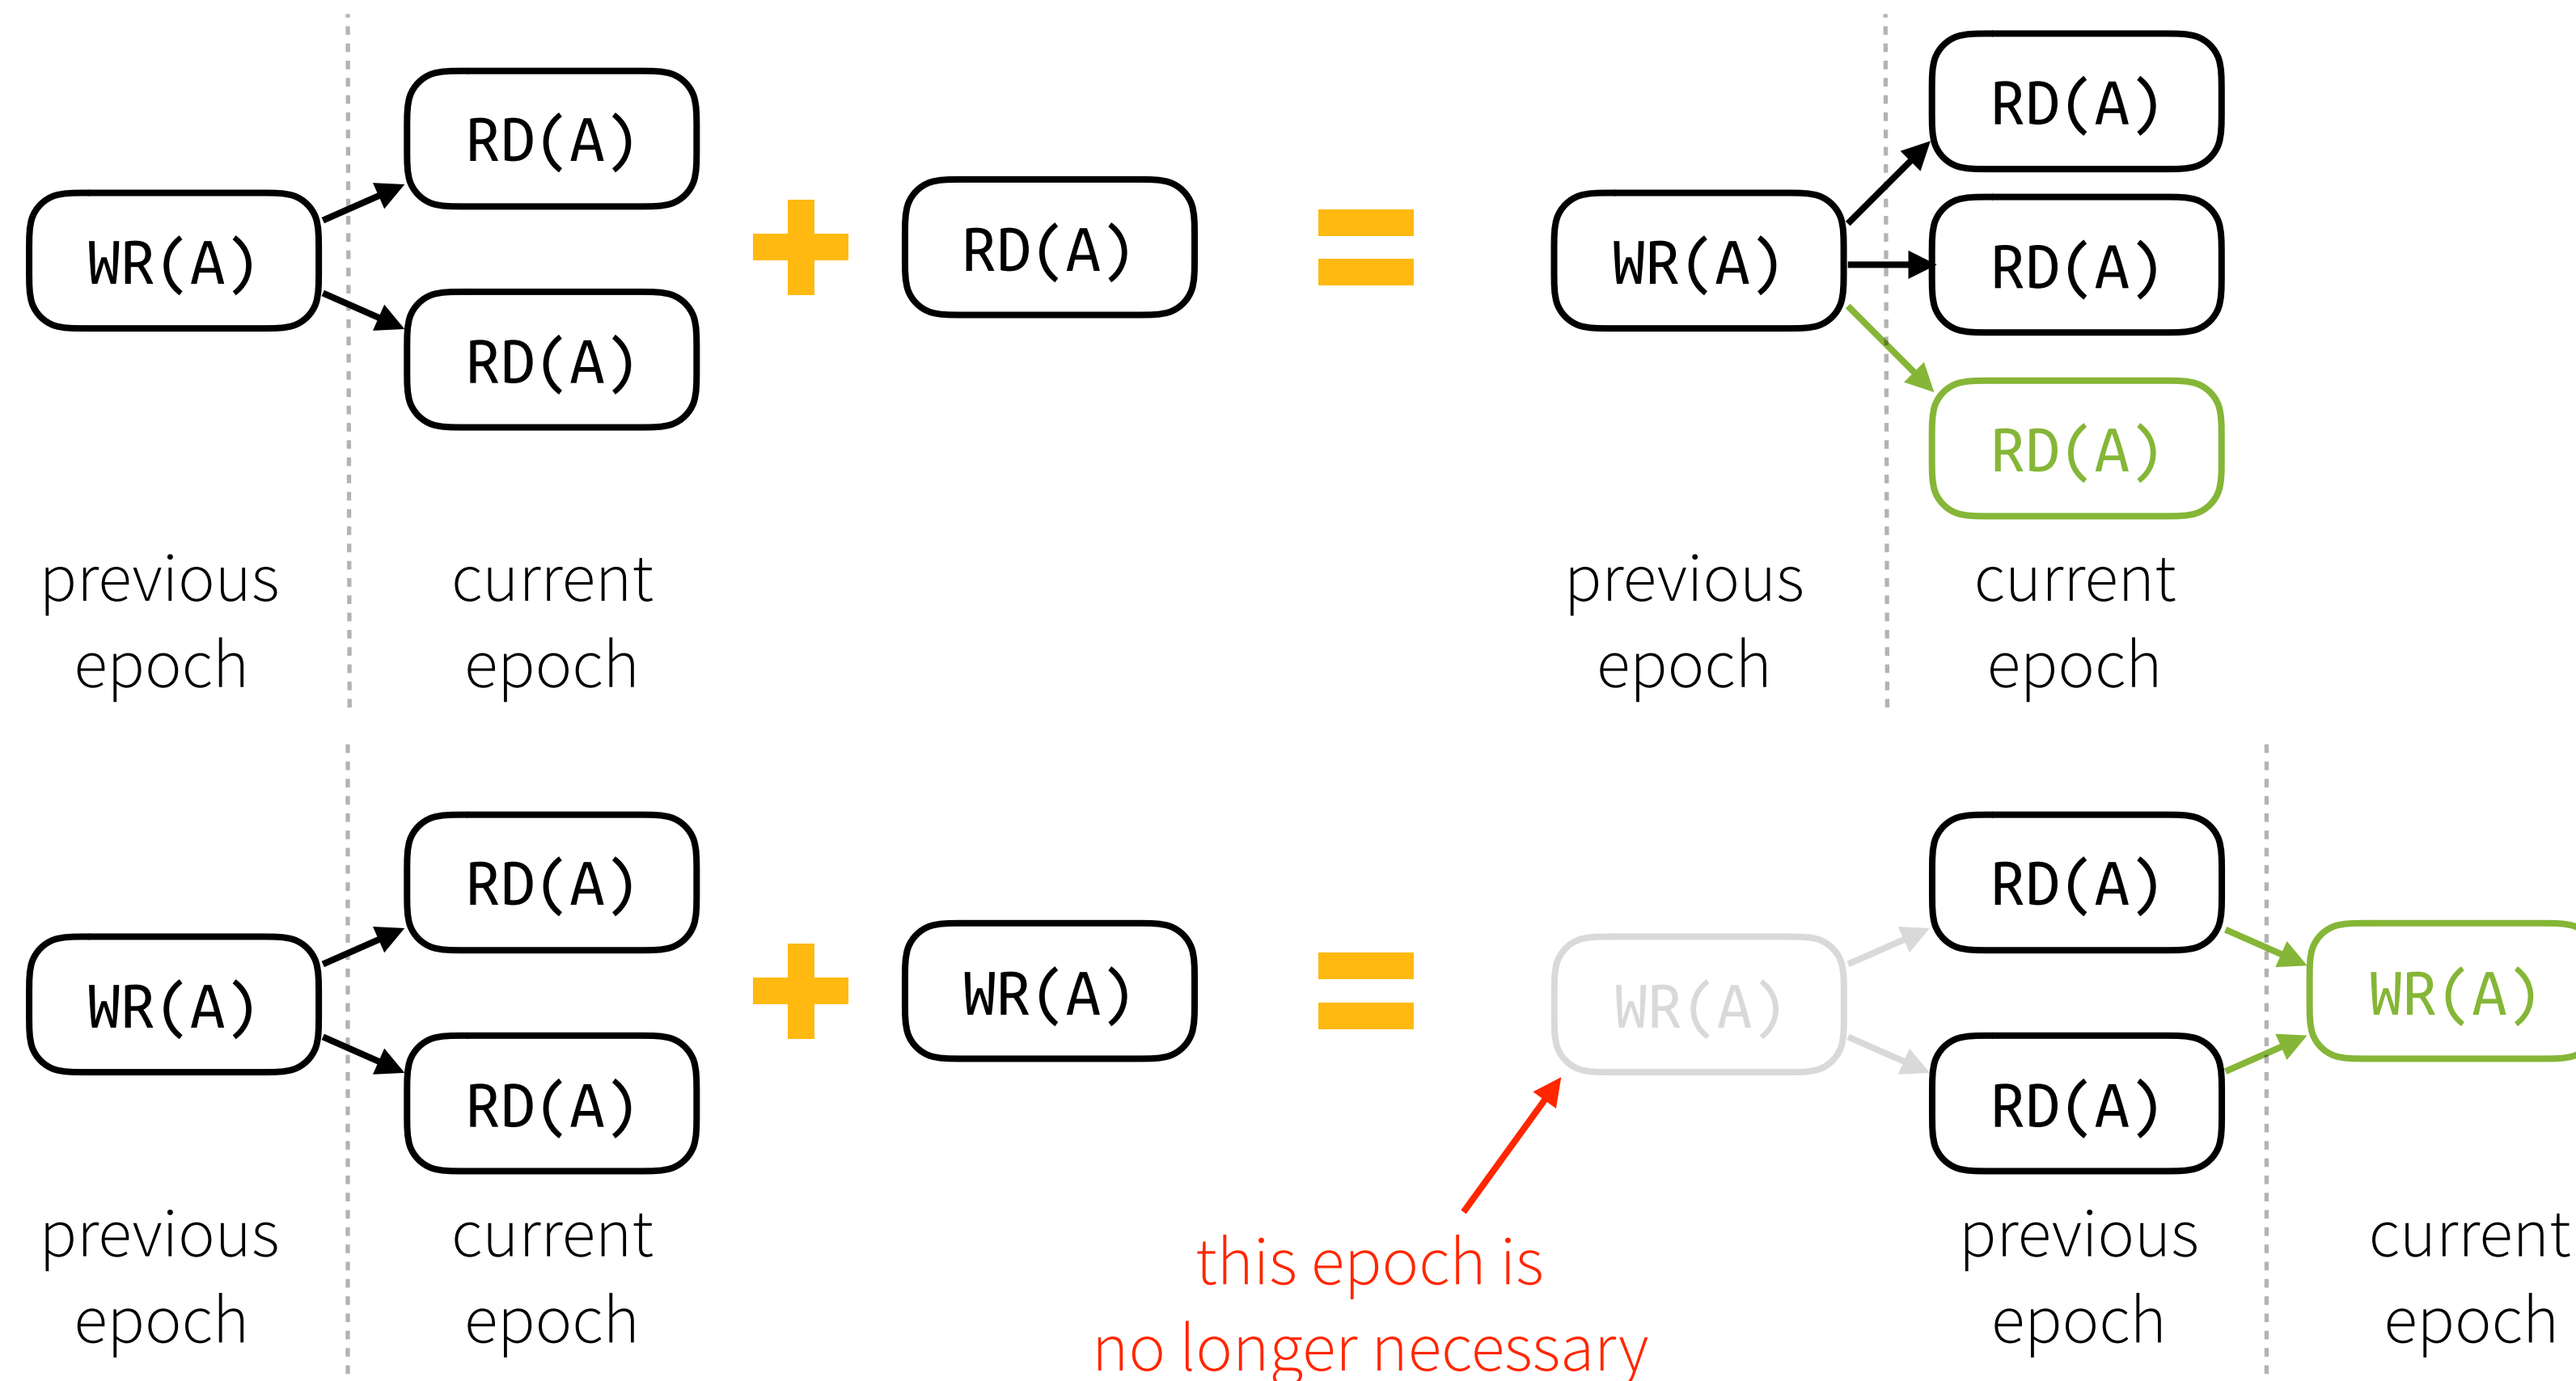
this dependence is **transitively** expressed by in-between epochs



- Only the dependencies between immediate epochs are non-transitive

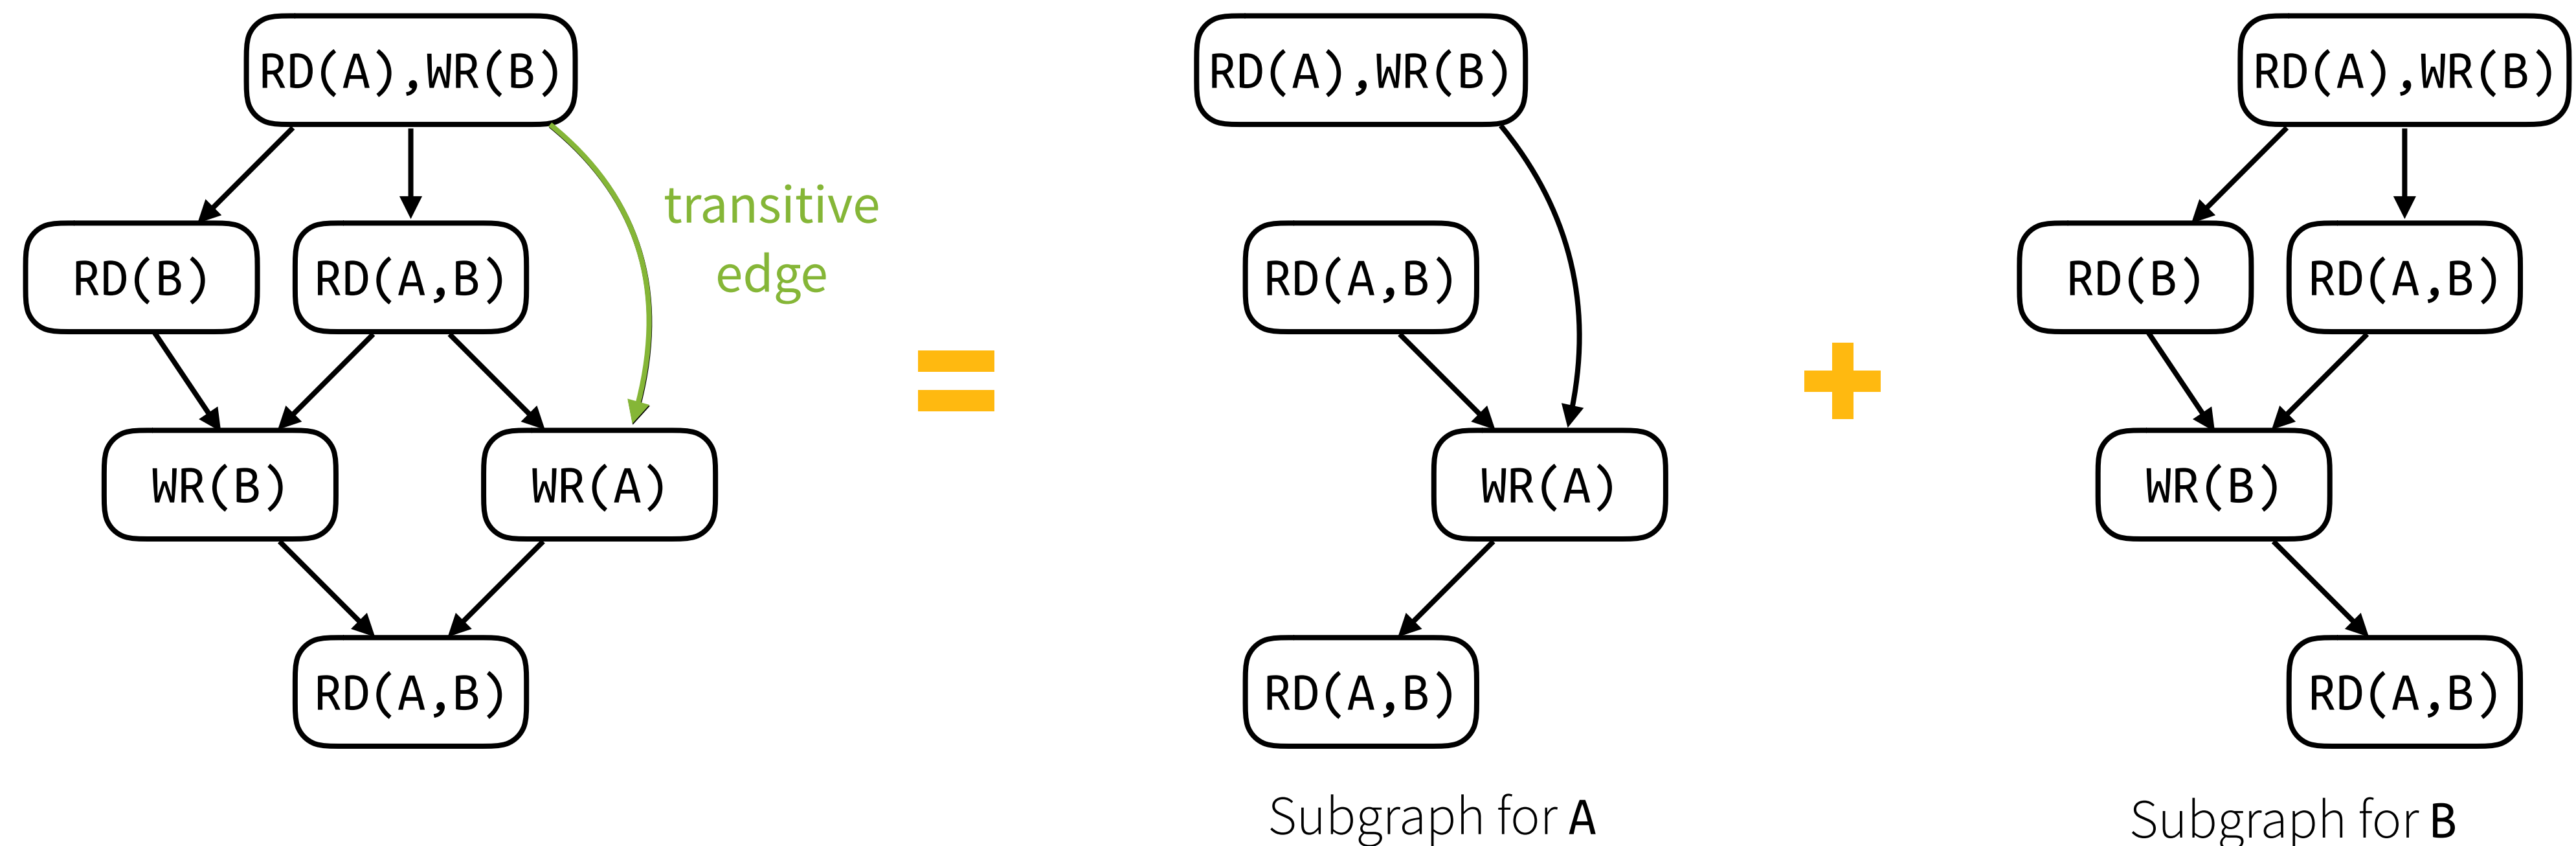
D_{epoch} : Epoch-based Algorithm

- Based on SWMR invariant: either **S**ingle **W**riter or **M**ultiple **R**eaders can exist for a region
- “Two-epoch rule”: remembering only the last two epochs is enough to find all dependencies



D_{epoch} : Epoch-based Algorithm

- Key idea: a task graph is decomposed into subgraphs for individual regions

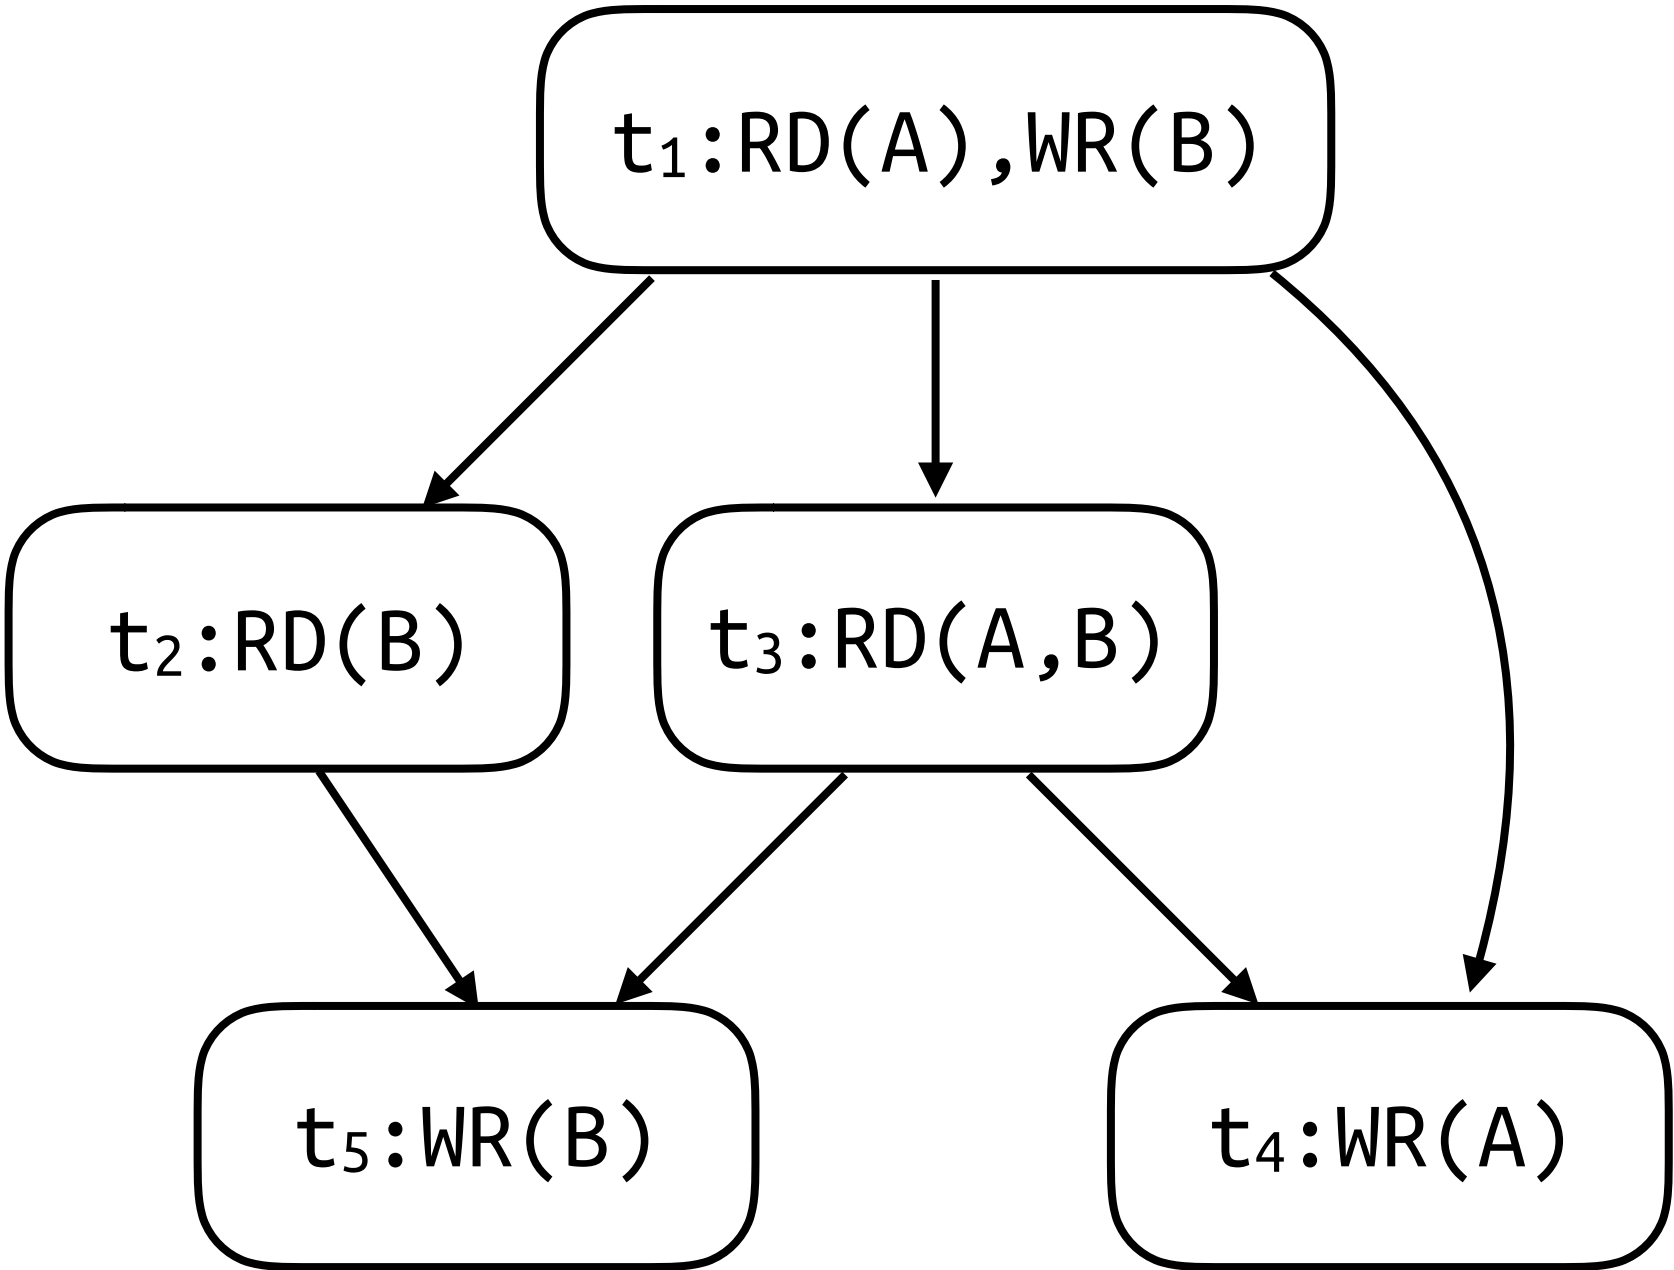


- Track epochs individually

- Each subgraph do not have transitive dependencies, but the union **might**

Example

Task graph



Task

- t₁: RD(A),WR(B)
- t₂: RD(B)
- t₃: RD(A,B)
- t₄: WR(A)
- t₅: WR(B)

Region A

previous epoch current epoch

Region B

previous epoch current epoch

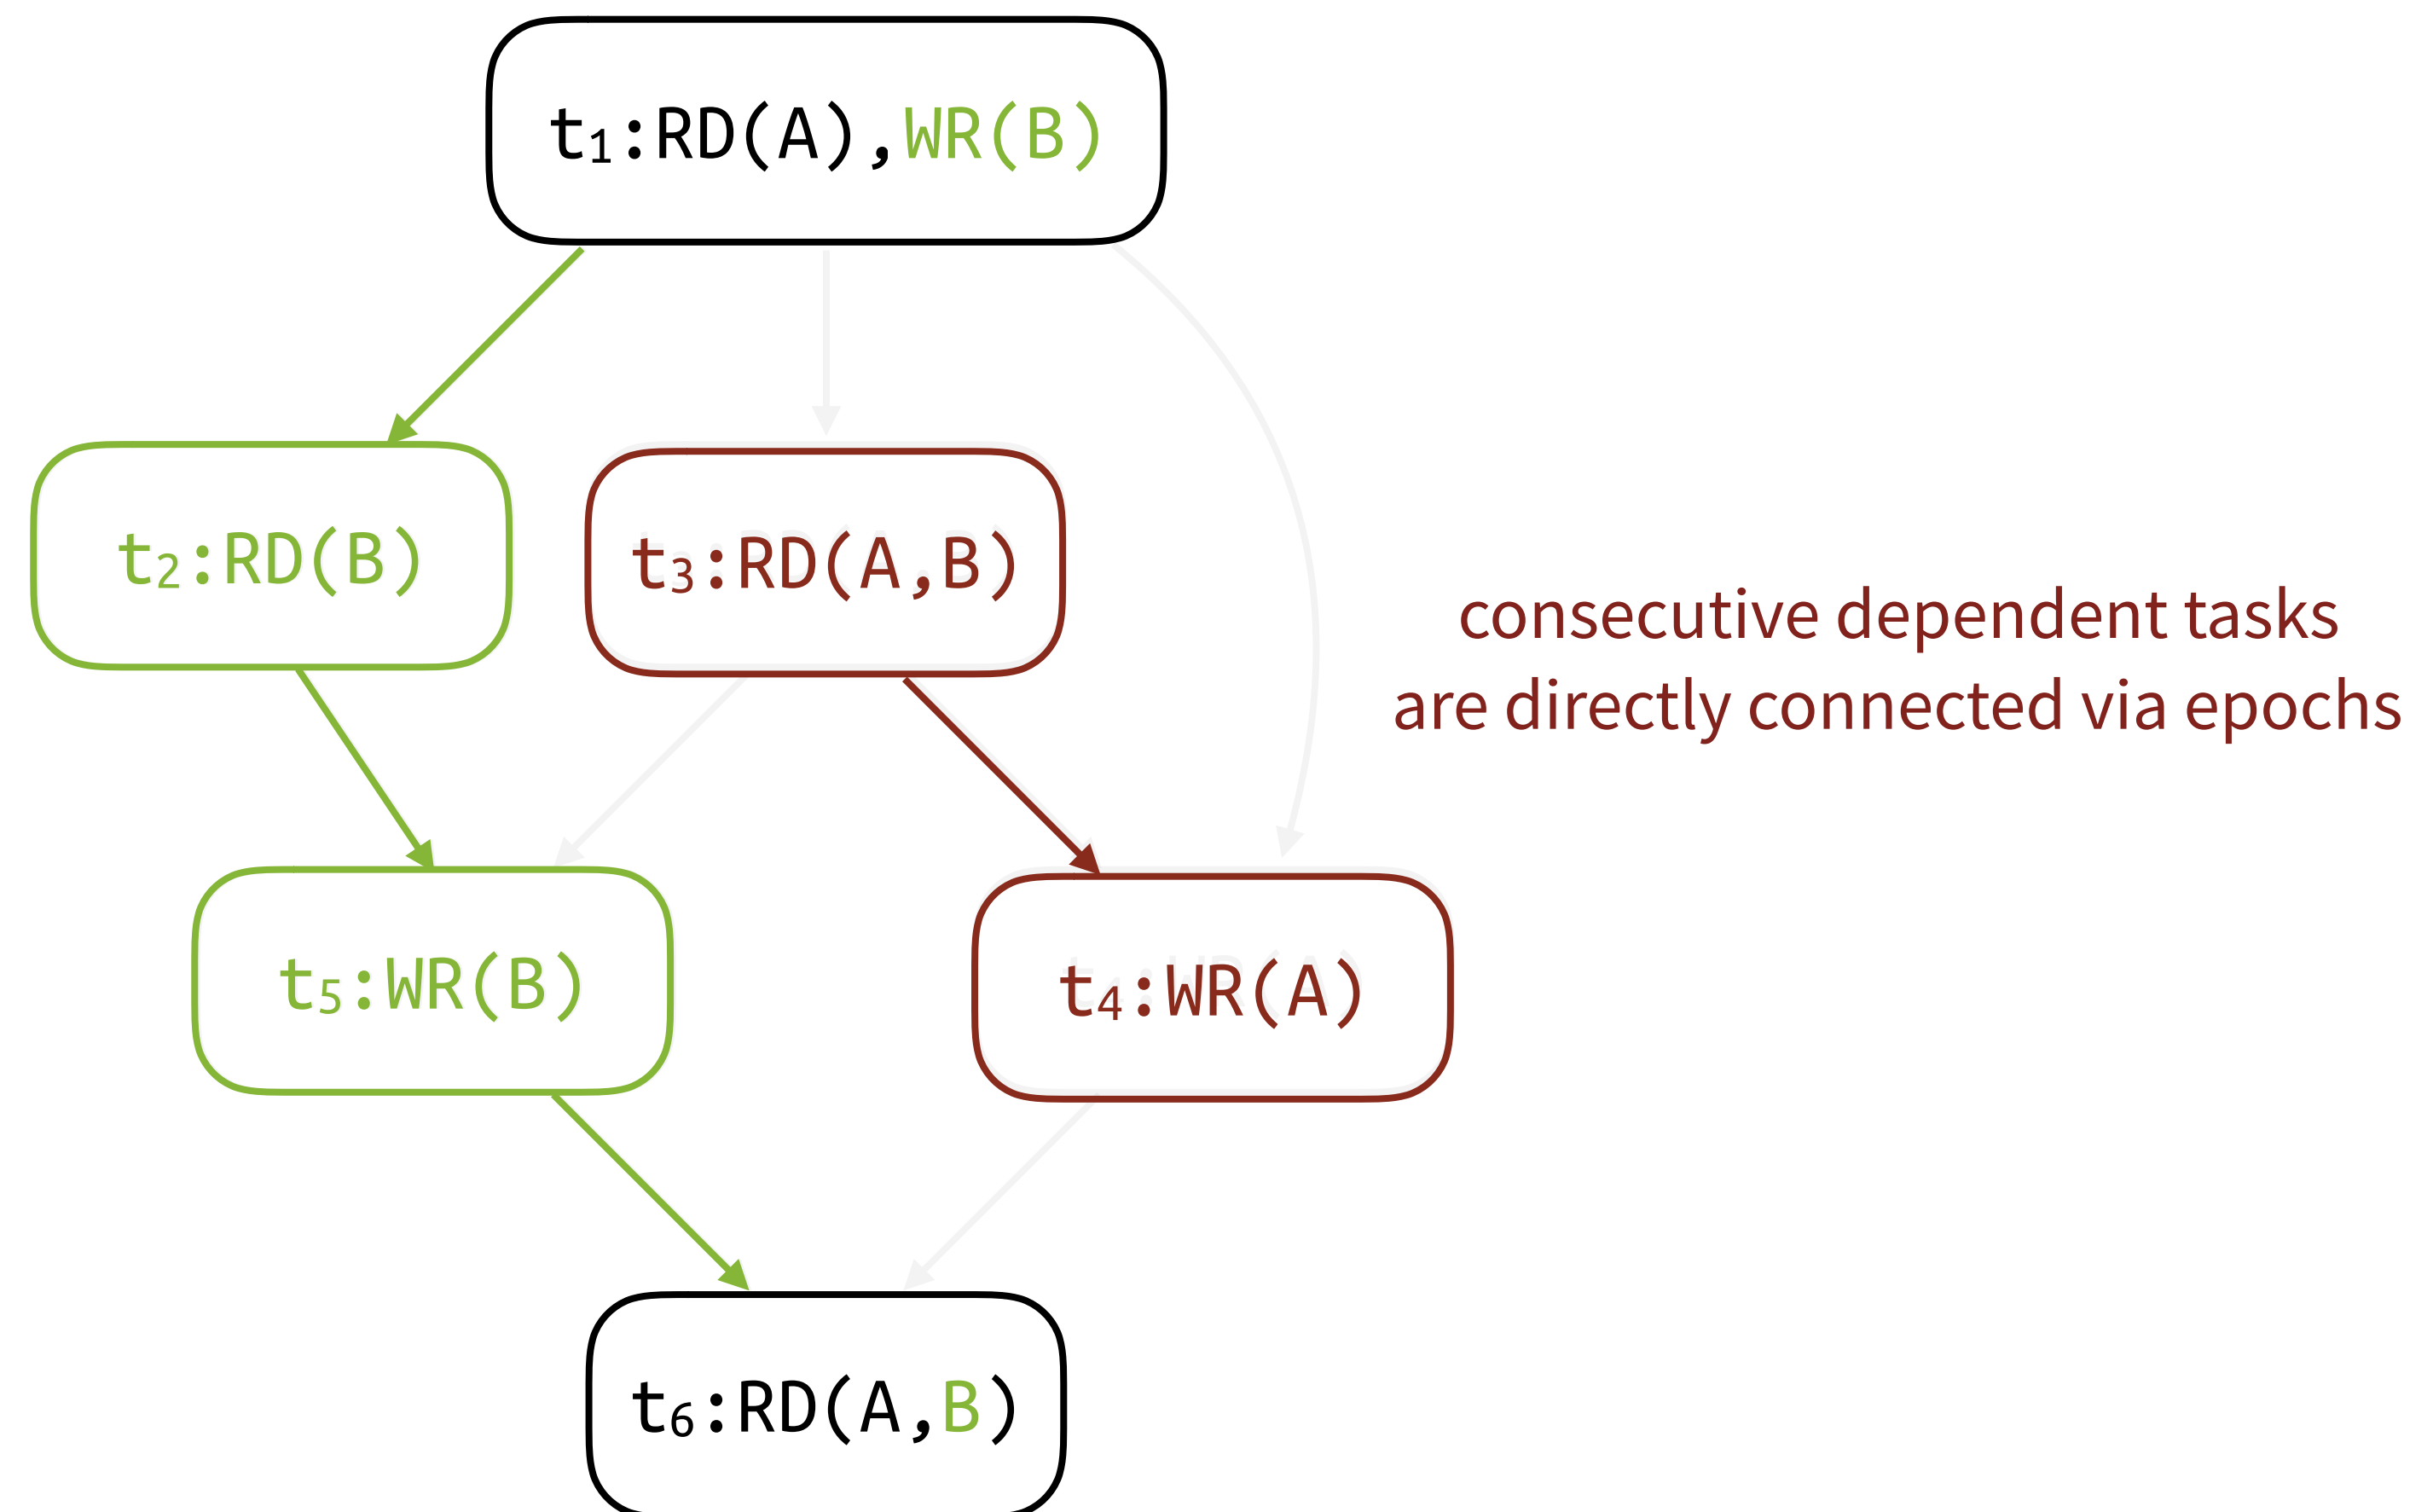
Region A		Region B	
previous epoch	current epoch	previous epoch	current epoch
	t ₁		t ₁
	t ₁	t ₁	t ₂
	t ₁ , t ₃	t ₁	t ₂ , t ₃
t ₁ , t ₃	t ₄	t ₁	t ₂ , t ₃
t ₁ , t ₃	t ₄	t ₂ , t ₃	t ₅

Soundness and Completeness of D_{epoch}

- By showing D_{epoch} is equivalent to a naive method D_{simple}
 - $D_{\text{simple}}(t_1, \dots, t_n) = \text{connect every two tasks } t_i \text{ and } t_j \text{ (} i < j \text{) when } t_i \text{ and } t_j \text{ are dependent}$
 - D_{simple} is sound and complete
- For every edge in D_{simple} , we show there is a path in D_{epoch} and vice versa

Soundness and Completeness of D_{epoch}

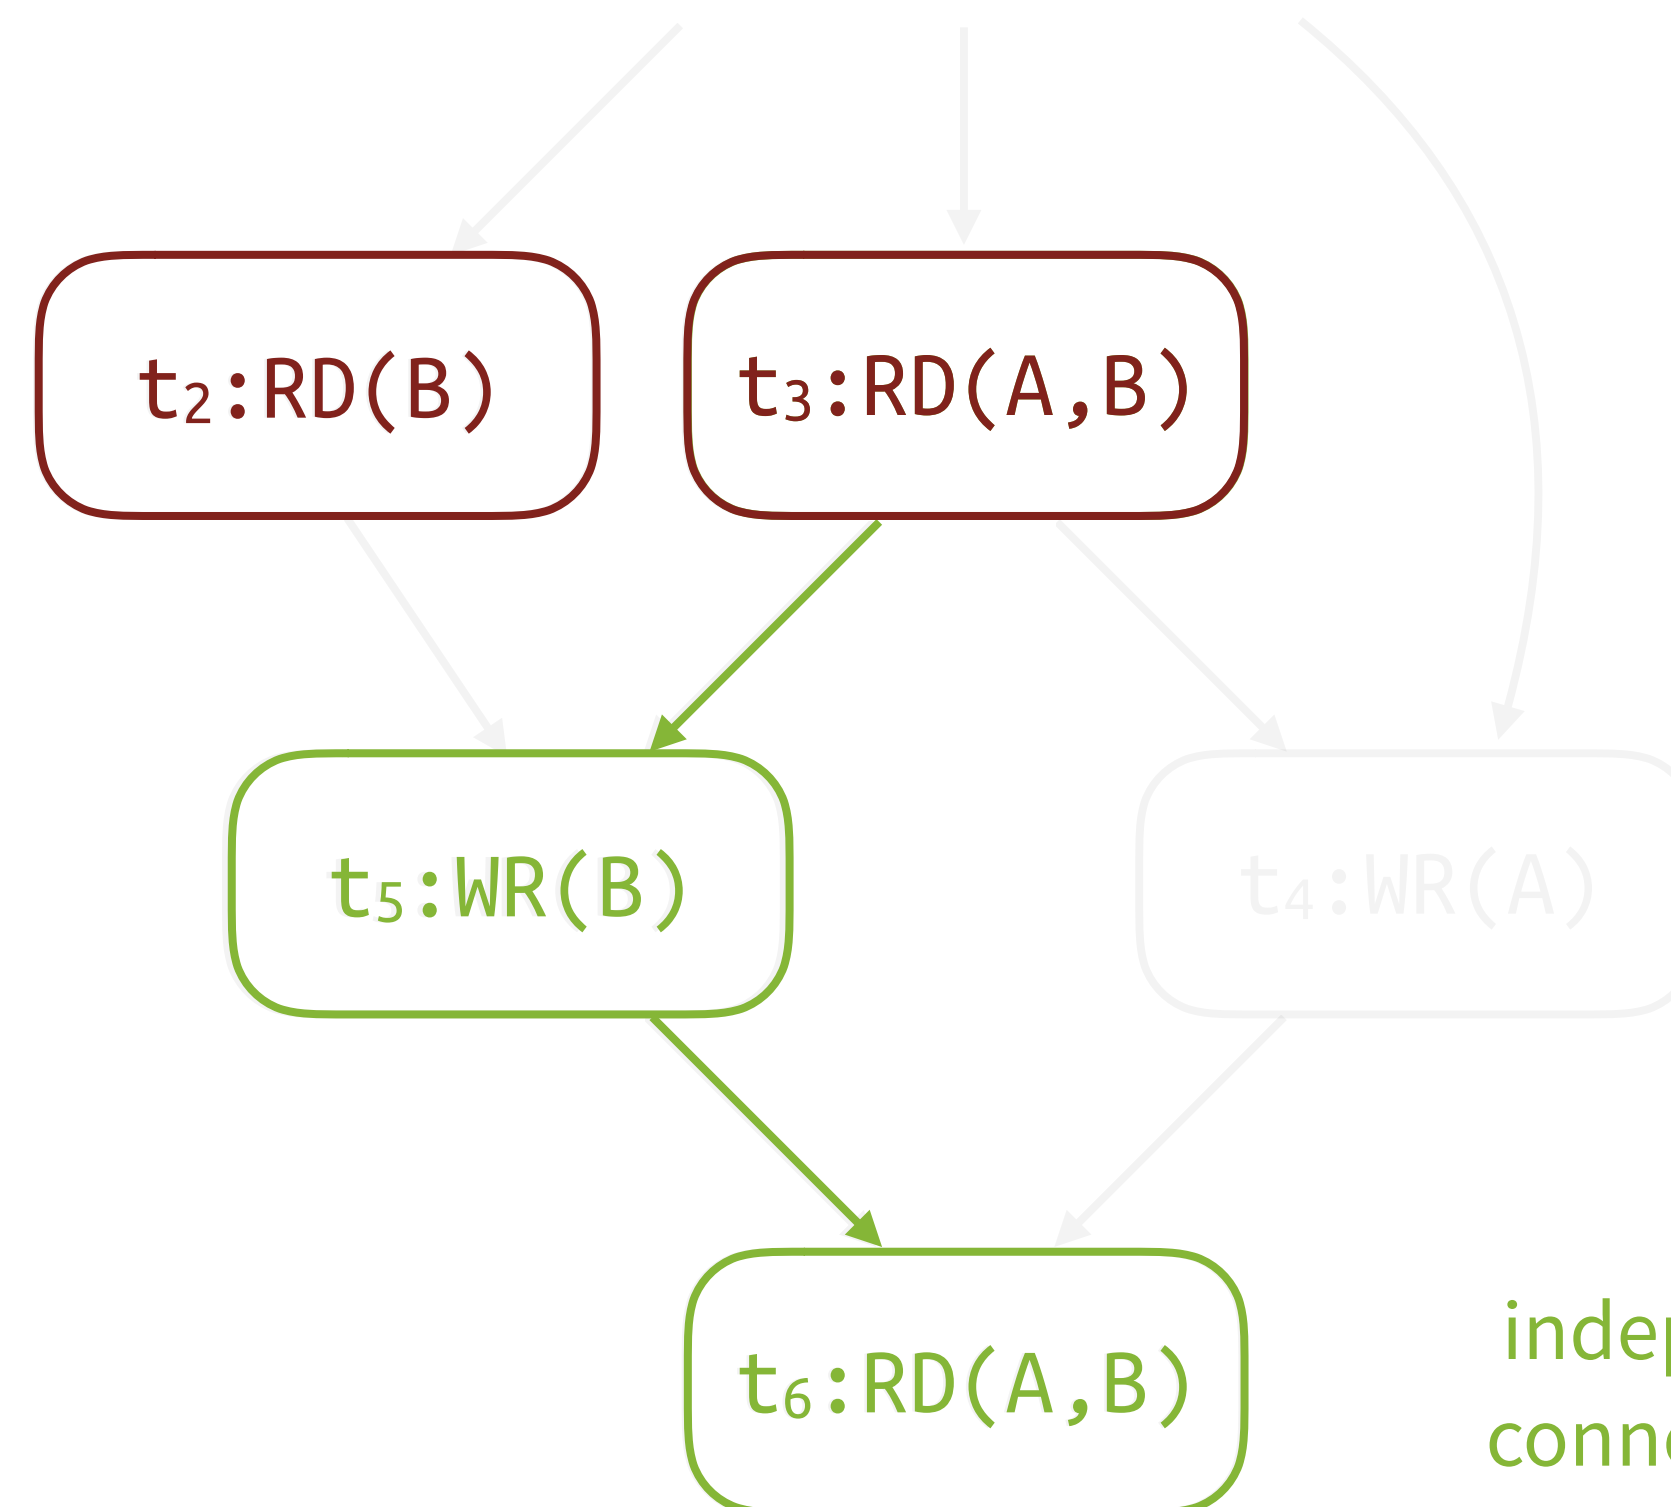
- Any two dependent tasks are **connected by at least one path**



Soundness and Completeness of D_{epoch}

- Independent tasks are never directly connected
- The length of every path between independent tasks is always greater than 1

consecutive independent tasks
must be mutually unreachable

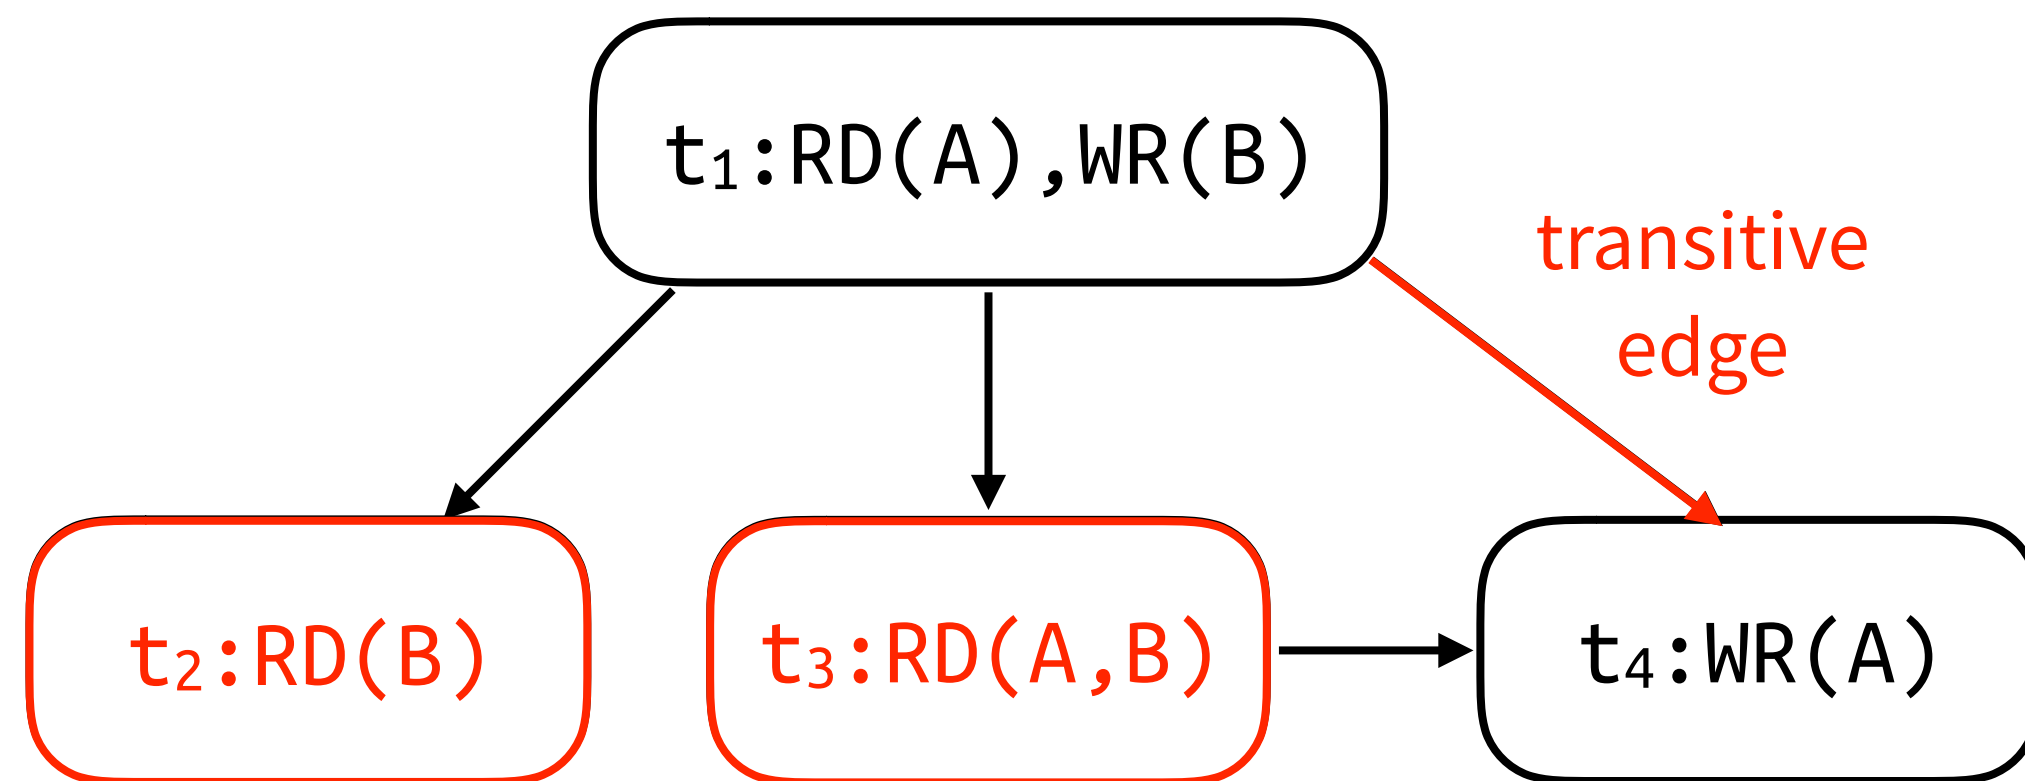


independent tasks are always
connected via some other task

When D_{epoch} Produces Parsimonious Task Graphs

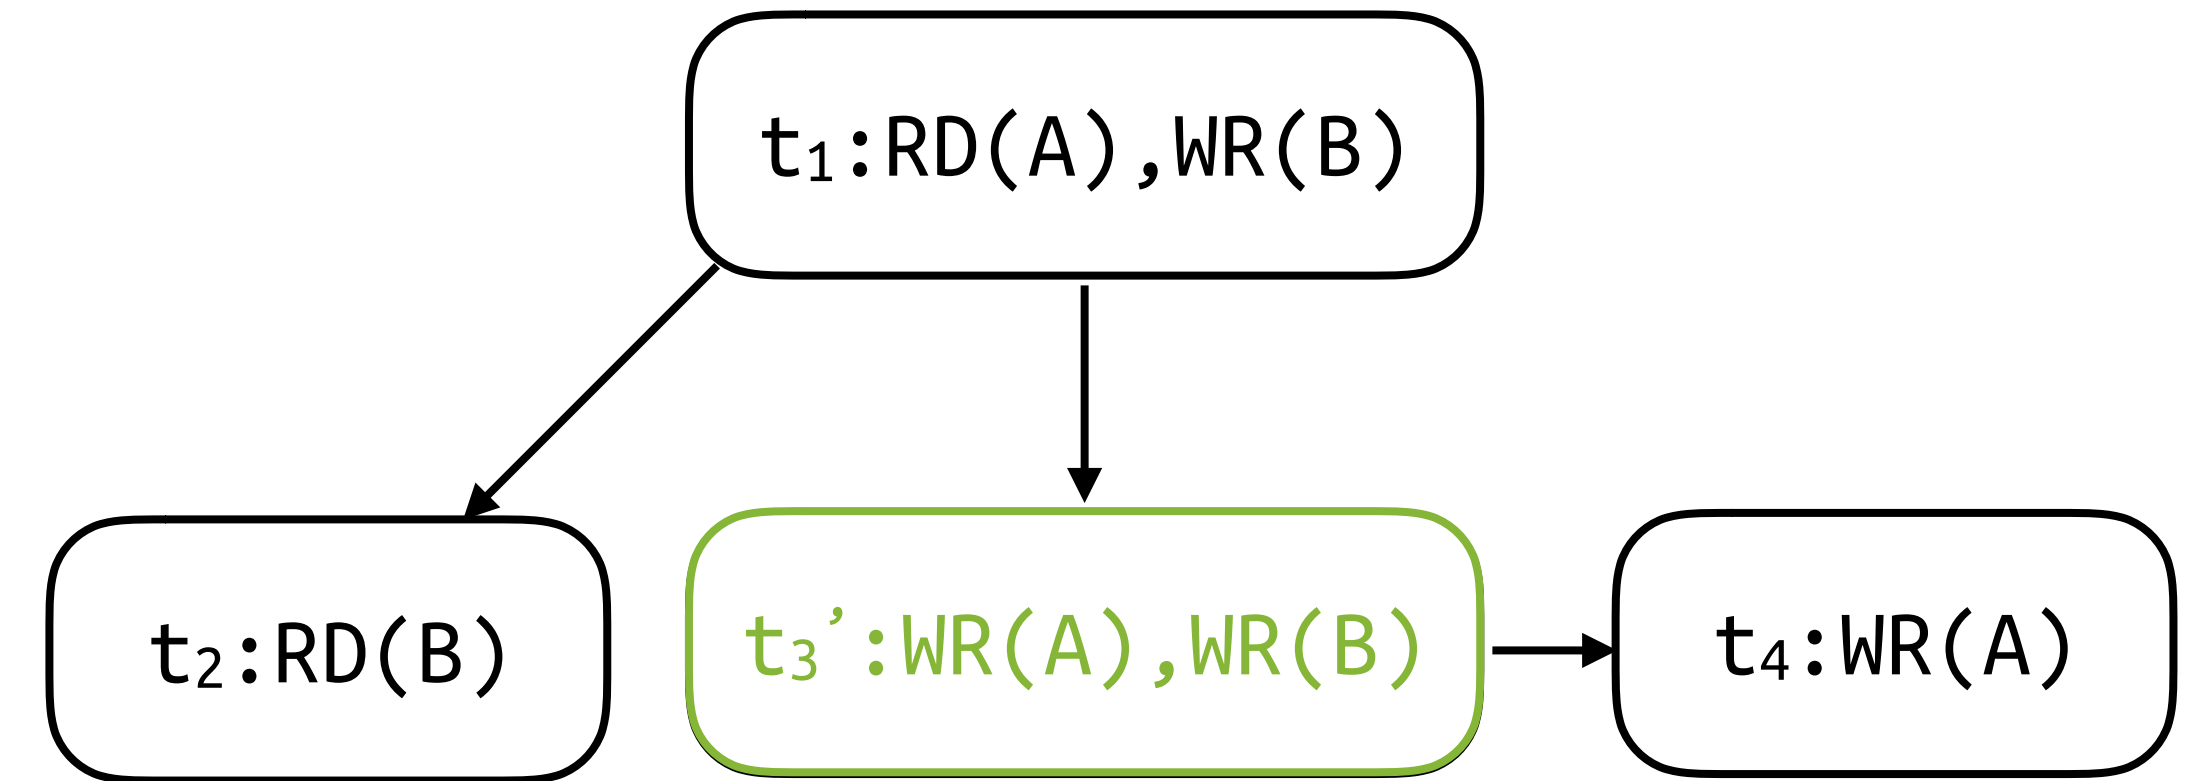
- When no “triangle” exists in a task stream

Non-parsimonious graph



because no in-between tasks
“redefine” A

Parsimonious graph



t_3' “cuts” dependence $t_1 \rightarrow t_4$

- For any subsequence t_1, \dots, t_n , if t_1 and t_n are dependent because of R , there must be a t_i that writes R

Conclusion

- Epoch-based dependence analysis algorithm D_{epoch} is sound and complete
 - And sometimes produces parsimonious task graphs
 - Even when the task graph is concurrently mutated for execution (see the paper)

Future Work

- Add advanced features in Legion to D_{epoch}
 - Coherence: multiple instances per region
 - Aliasing: regions can be aliased
 - Tracing: task graphs can be memoized
 - Resilience: failed tasks can be recovered
 - Replication: dependence analysis can be replicated

Acknowledgment

This research was supported by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration, award DE-NA0002373-1 from the Department of Energy National Nuclear Security Administration, NSF grant CCF-1160904, and an internship at Los Alamos National Laboratory.

Questions?