

# FRACTAL

AN EXECUTION MODEL FOR FINE-GRAIN  
NESTED SPECULATIVE PARALLELISM

---

SUVINAY SUBRAMANIAN, MARK C. JEFFREY,  
MALEEN ABEYDEERA, HYUN RYONG LEE, VICTOR A. YING,  
JOEL EMER, DANIEL SANCHEZ

ISCA 2017



**NVIDIA**

# Current speculative systems scale poorly

---

Speculative parallelization, e.g. TM, **simplifies parallel programming**

**Performs poorly** on real world applications...

...because applications comprise large atomic tasks

# Large atomic tasks limit performance

---

## Database Transaction

query X  
...  
...  
update Z  
...  
...  
query U  
...  
...  
update V

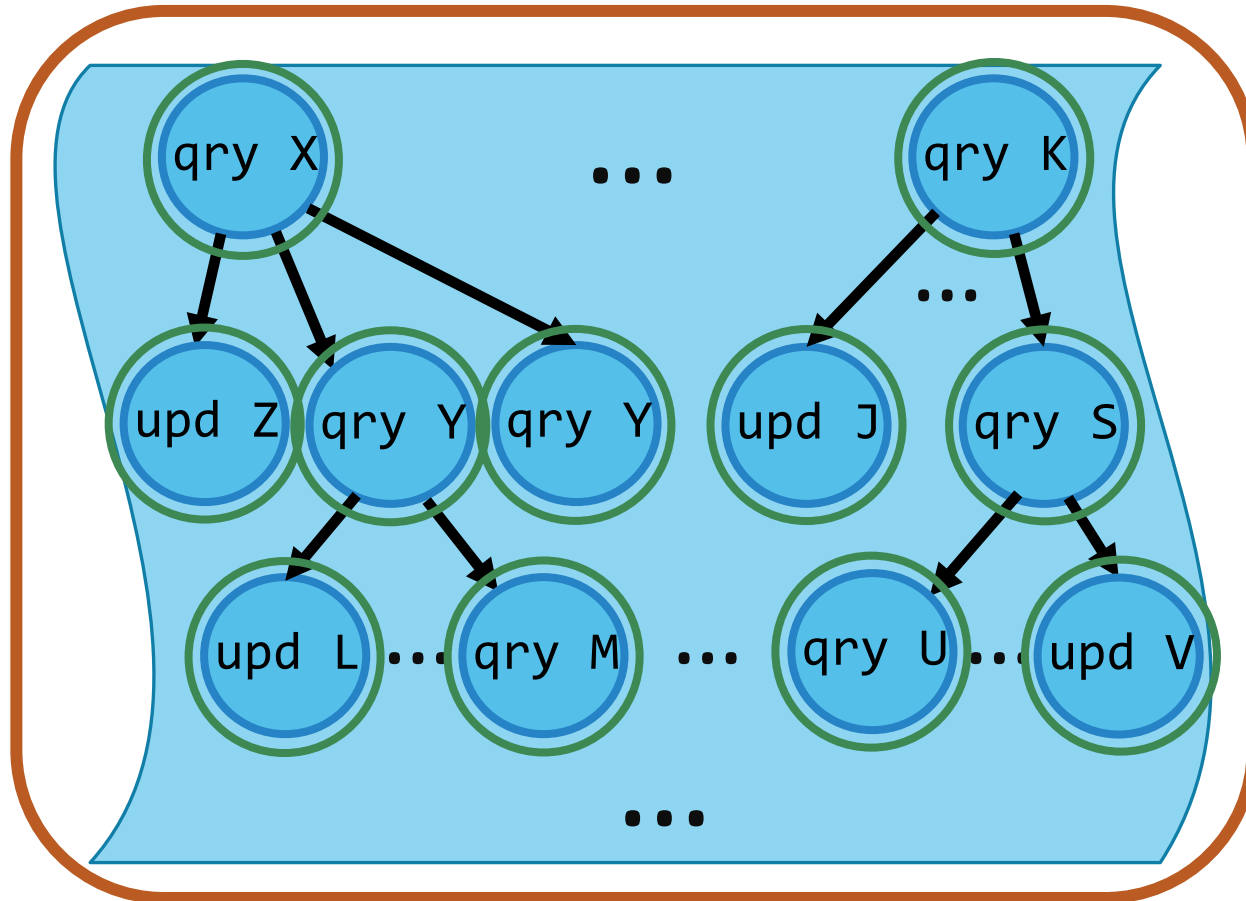
Millions of cycles

Prone to **aborts**

Challenging to track

**Serial** (misses parallelism)

# Large atomic tasks have abundant nested parallelism!

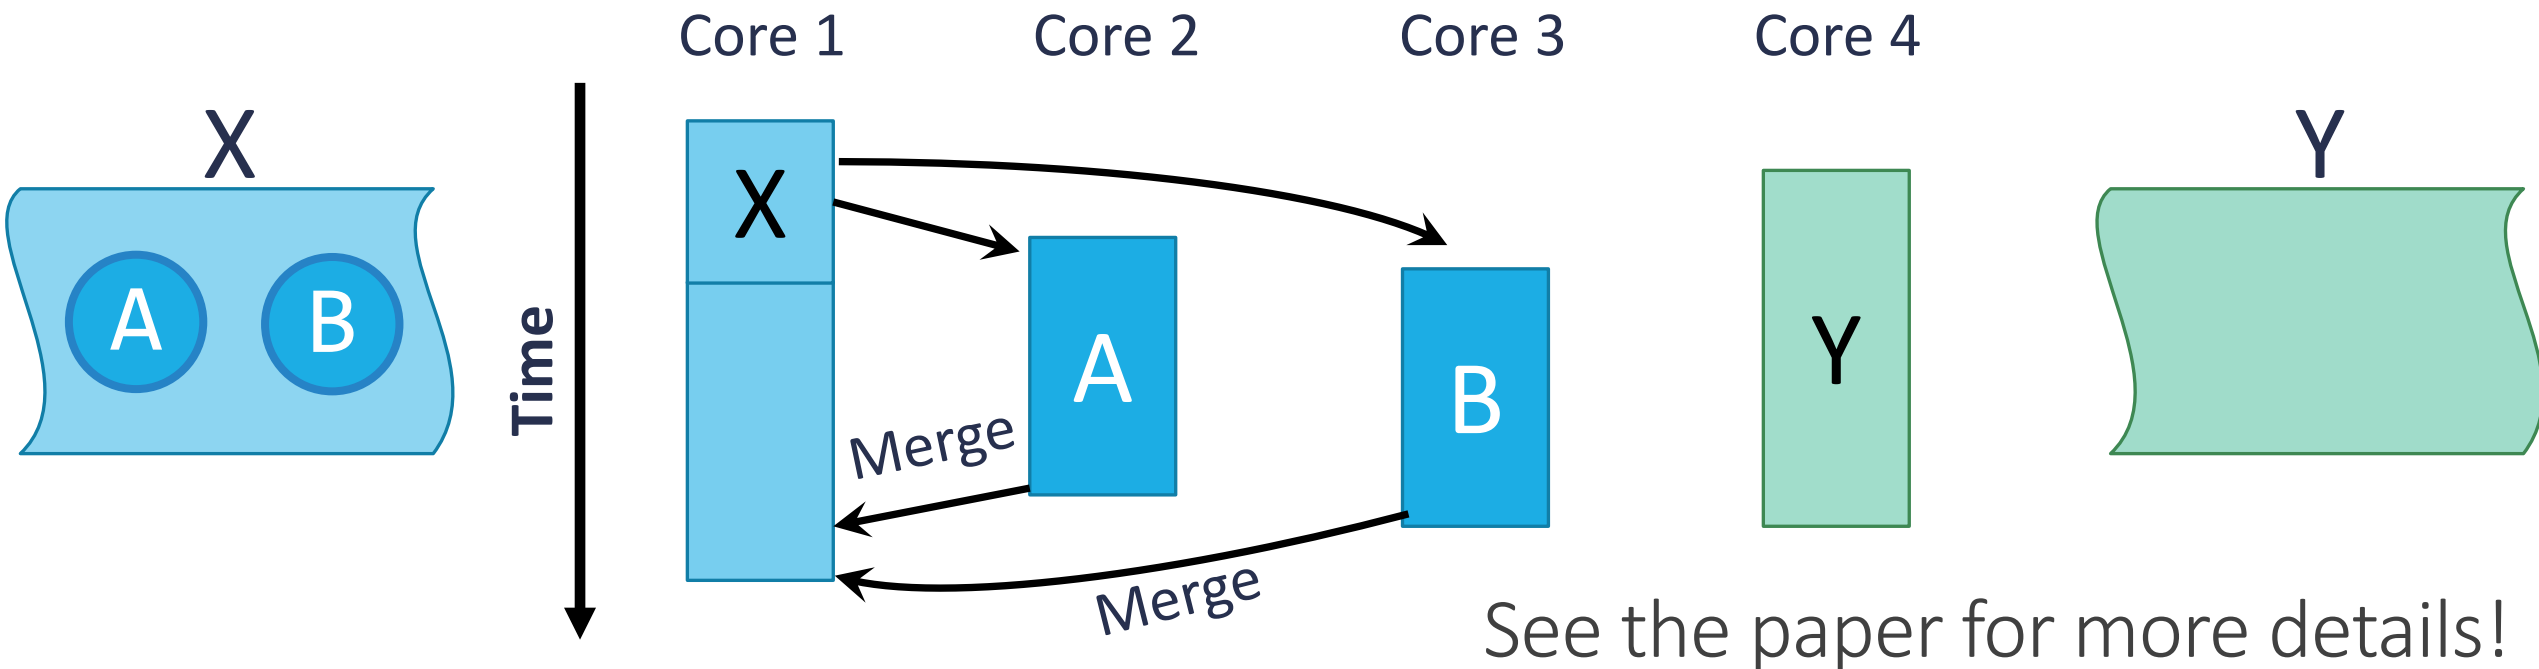


How to

- extract parallelism?
- maintain atomicity?
- achieve high performance?

# Prior TMs fail to exploit nested parallelism

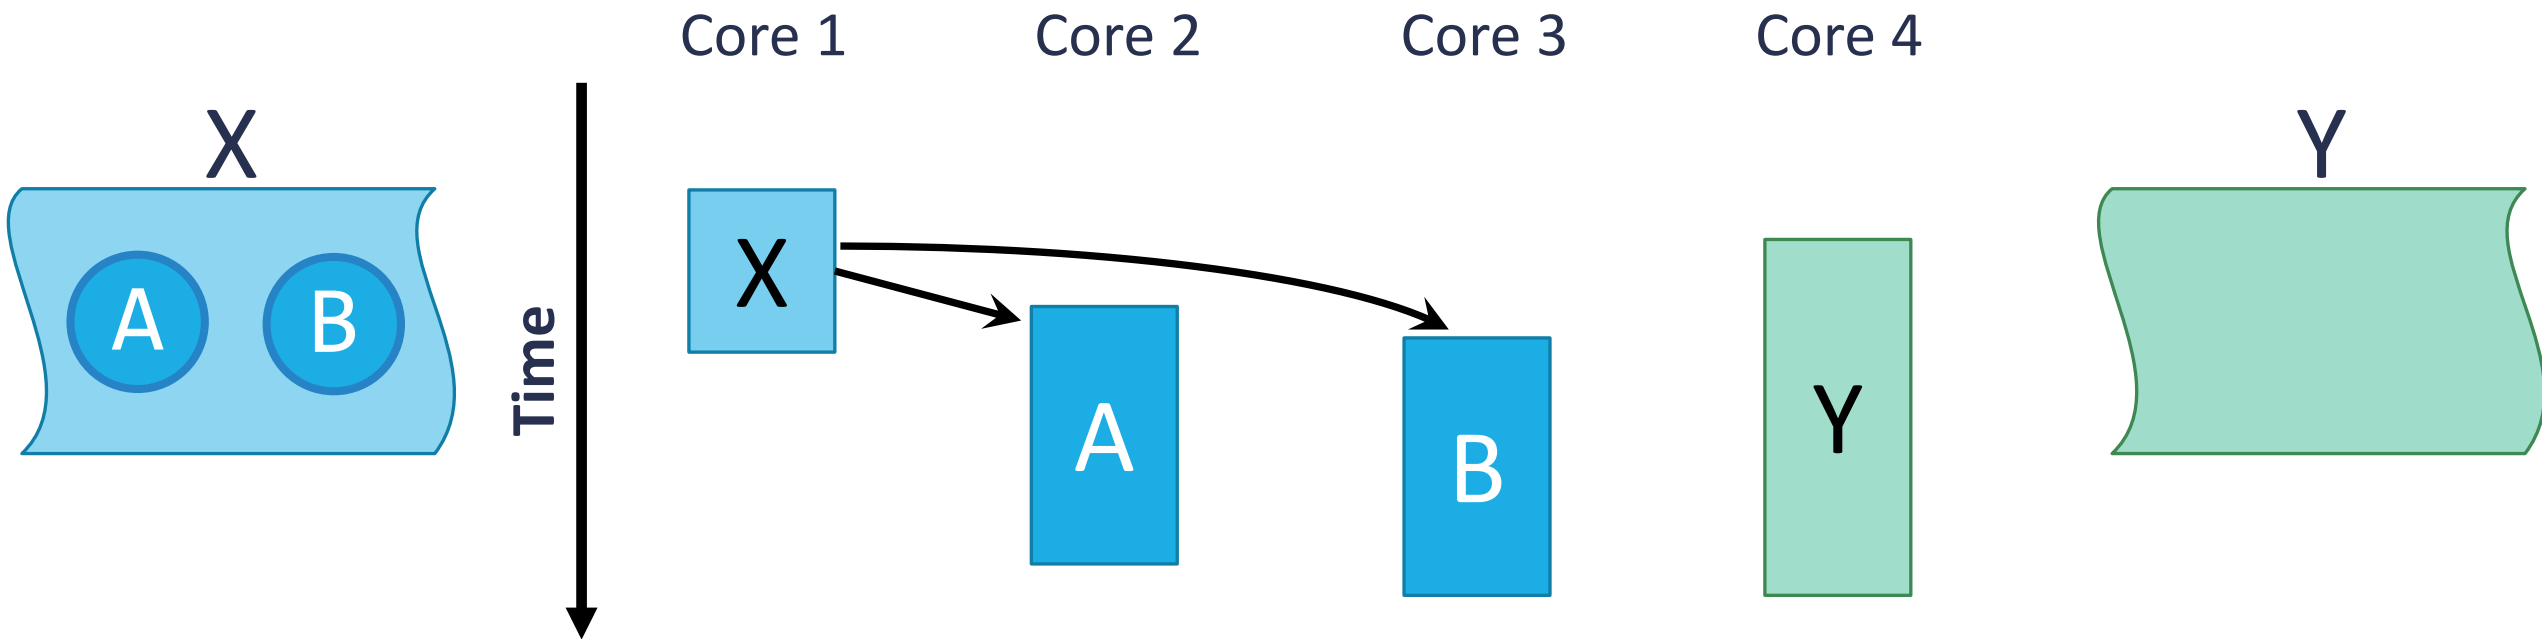
1. Merging of “nested” speculative state with parent  
Large speculative state, prone to aborts
2. Cyclic dependence between parent and nested children  
Deadlock and livelock issues



# Ordering tasks to guarantee atomicity

X  
1    Y  
2

X    A    B    Y  
1    1.1    1.2    2



# Fractal decouples atomicity from parallelism

---

1. Decouples unit of atomicity from unit of parallelism
  - *Domain*: All tasks belonging to a domain appear to execute atomically
2. Implementation guarantees atomicity by ordering tasks
  - No merging speculative state

Benefits of Fractal



Tiny tasks

Easy to track

Composable speculative parallelism

# Fractal Execution Model

---

DECOUPLING ATOMICITY FROM PARALLELISM




# *Domains* to group tasks into atomic units

Fractal programs consist of atomic tasks

Tasks may access arbitrary data

Tasks may create child tasks

Tasks belong to a hierarchy of nested domains



# Semantics across domains


Each task:

- can create a single *subdomain*
- can enqueue child tasks to *subdomain* or *current domain*

(All tasks in domain + creator of domain)



Appear to execute as single atomic unit



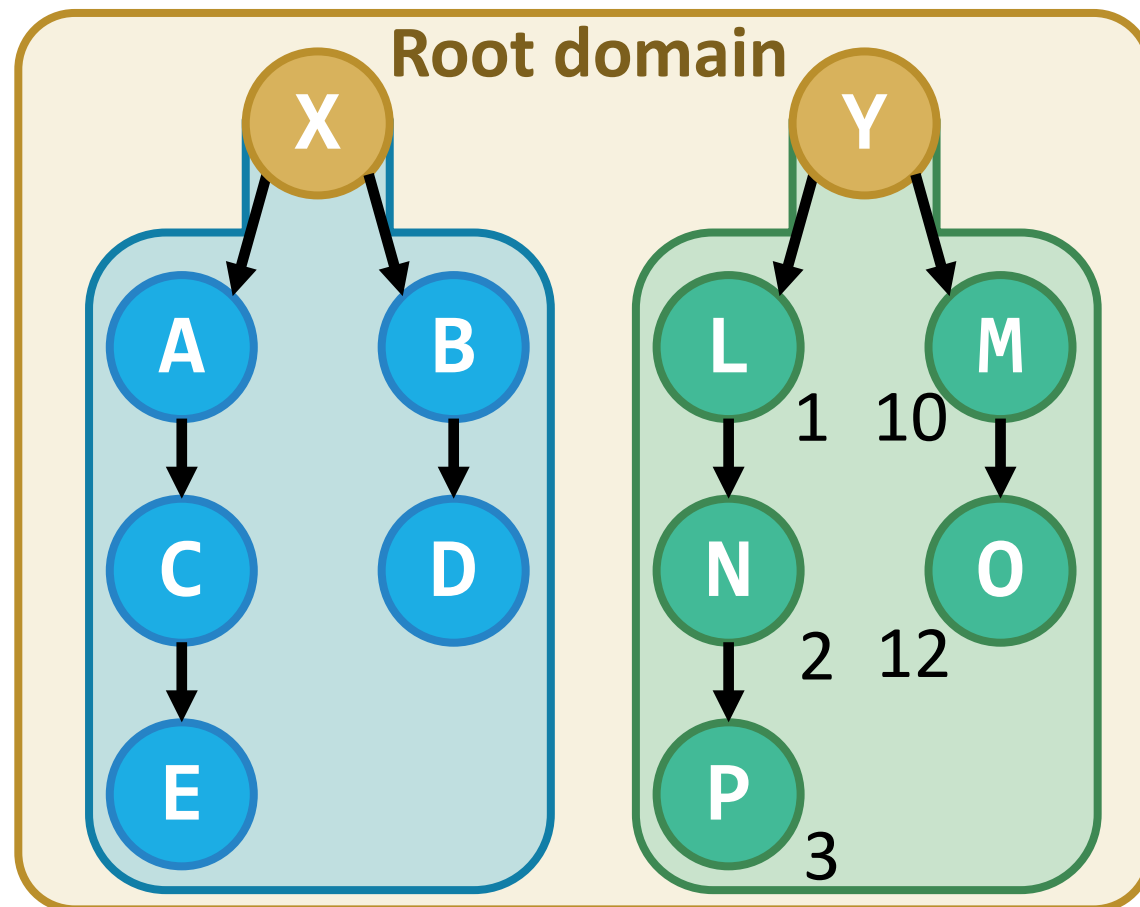
# Semantics within a domain

## Unordered

- Arbitrary order while respecting parent-child dependences

## Timestamp-ordered

- Tasks appear to execute in increasing timestamp order
- Children appear to execute after parent



# Fractal software API

---

## Creating and enqueueing tasks

```
fractal::enqueue(function_pointer,  
                 timestamp,  
                 arguments...);
```

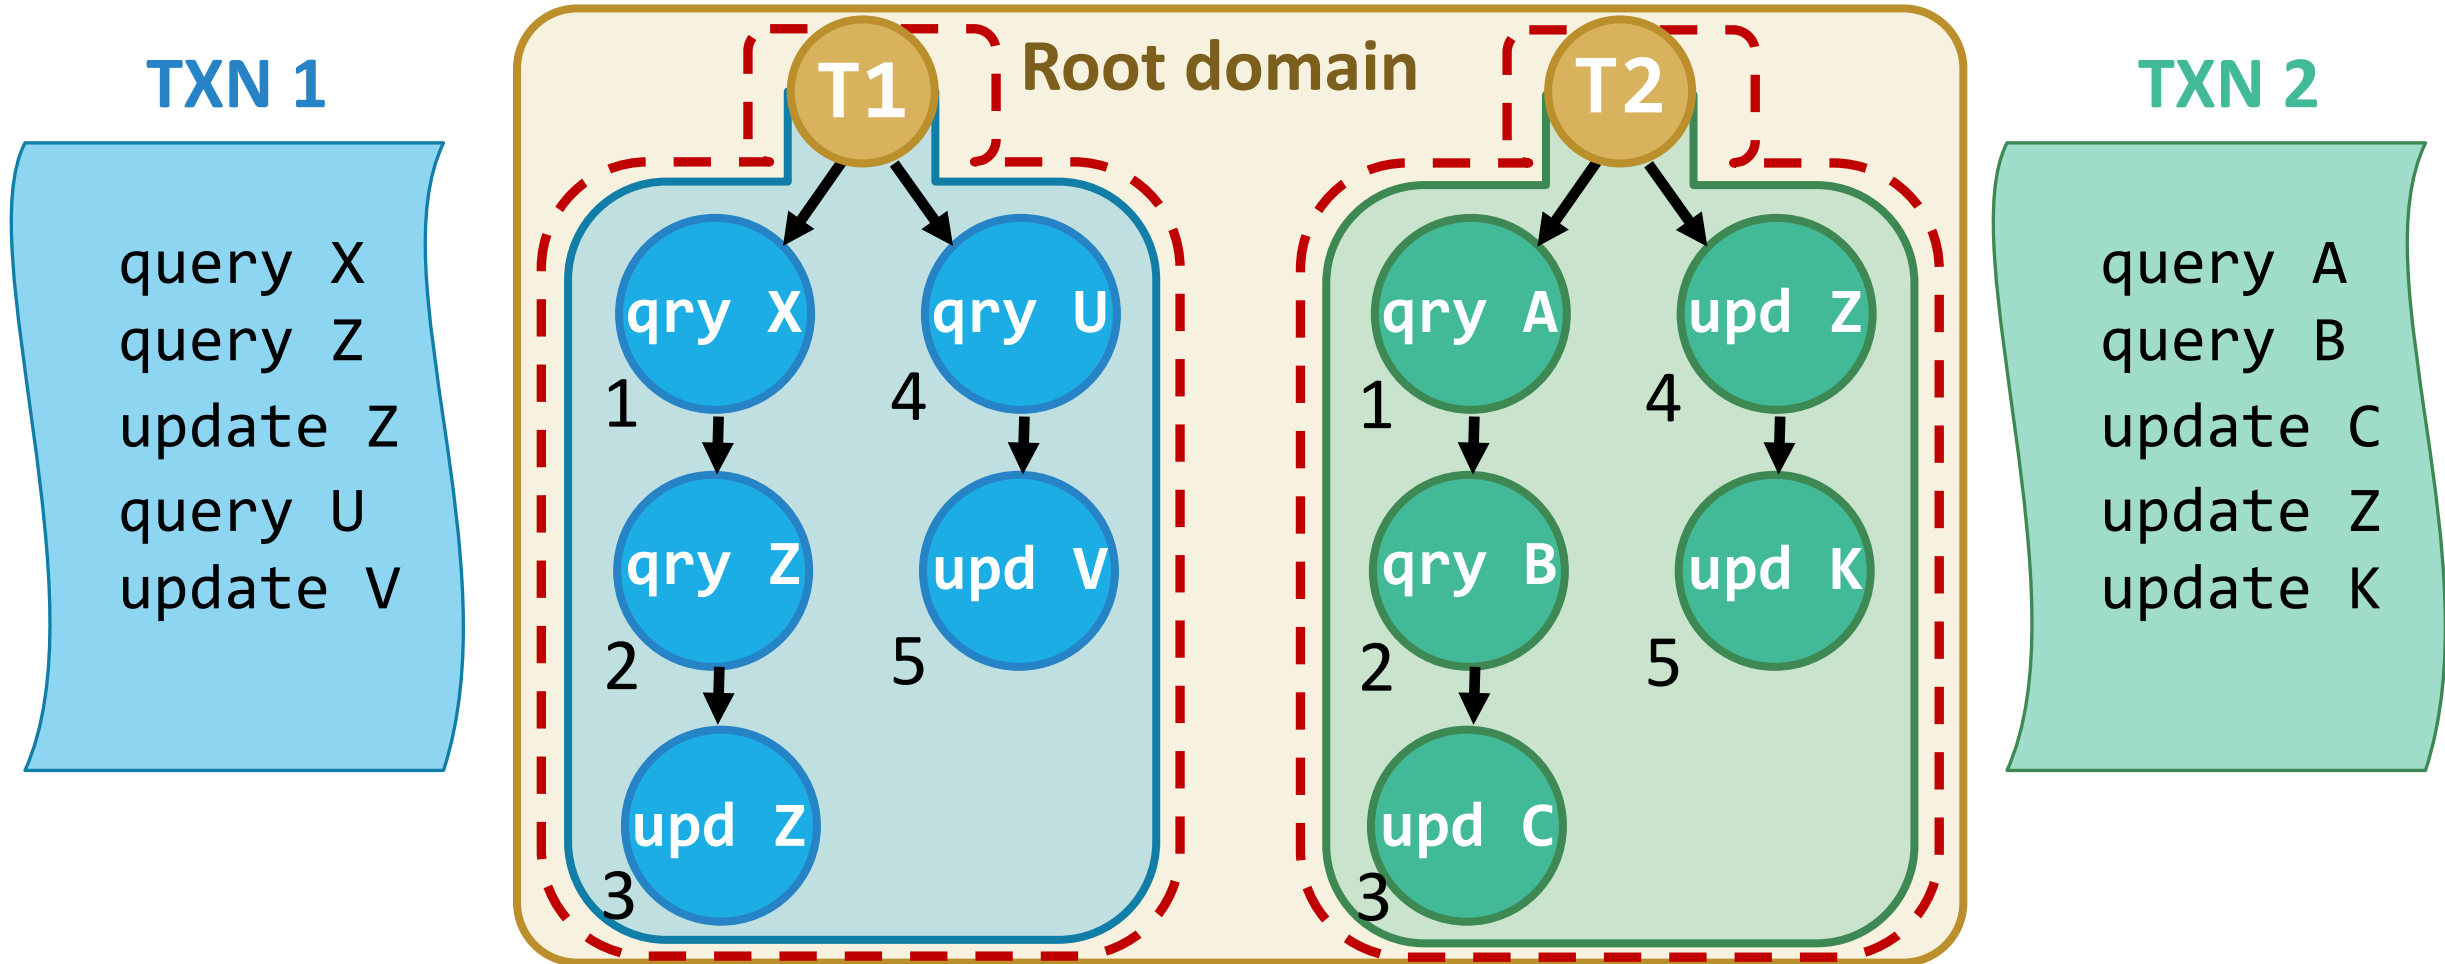
## Creating sub-domains

```
fractal::create_subdomain(<domain_type>);
```

## High-level programming interface, e.g.

```
forall(), callcc(), parallel_reduce()
```

# Example: Database transactions in Fractal



# Fractal Implementation

---


ATOMICITY THROUGH ORDERING

# Fractal Virtual Time (VT)

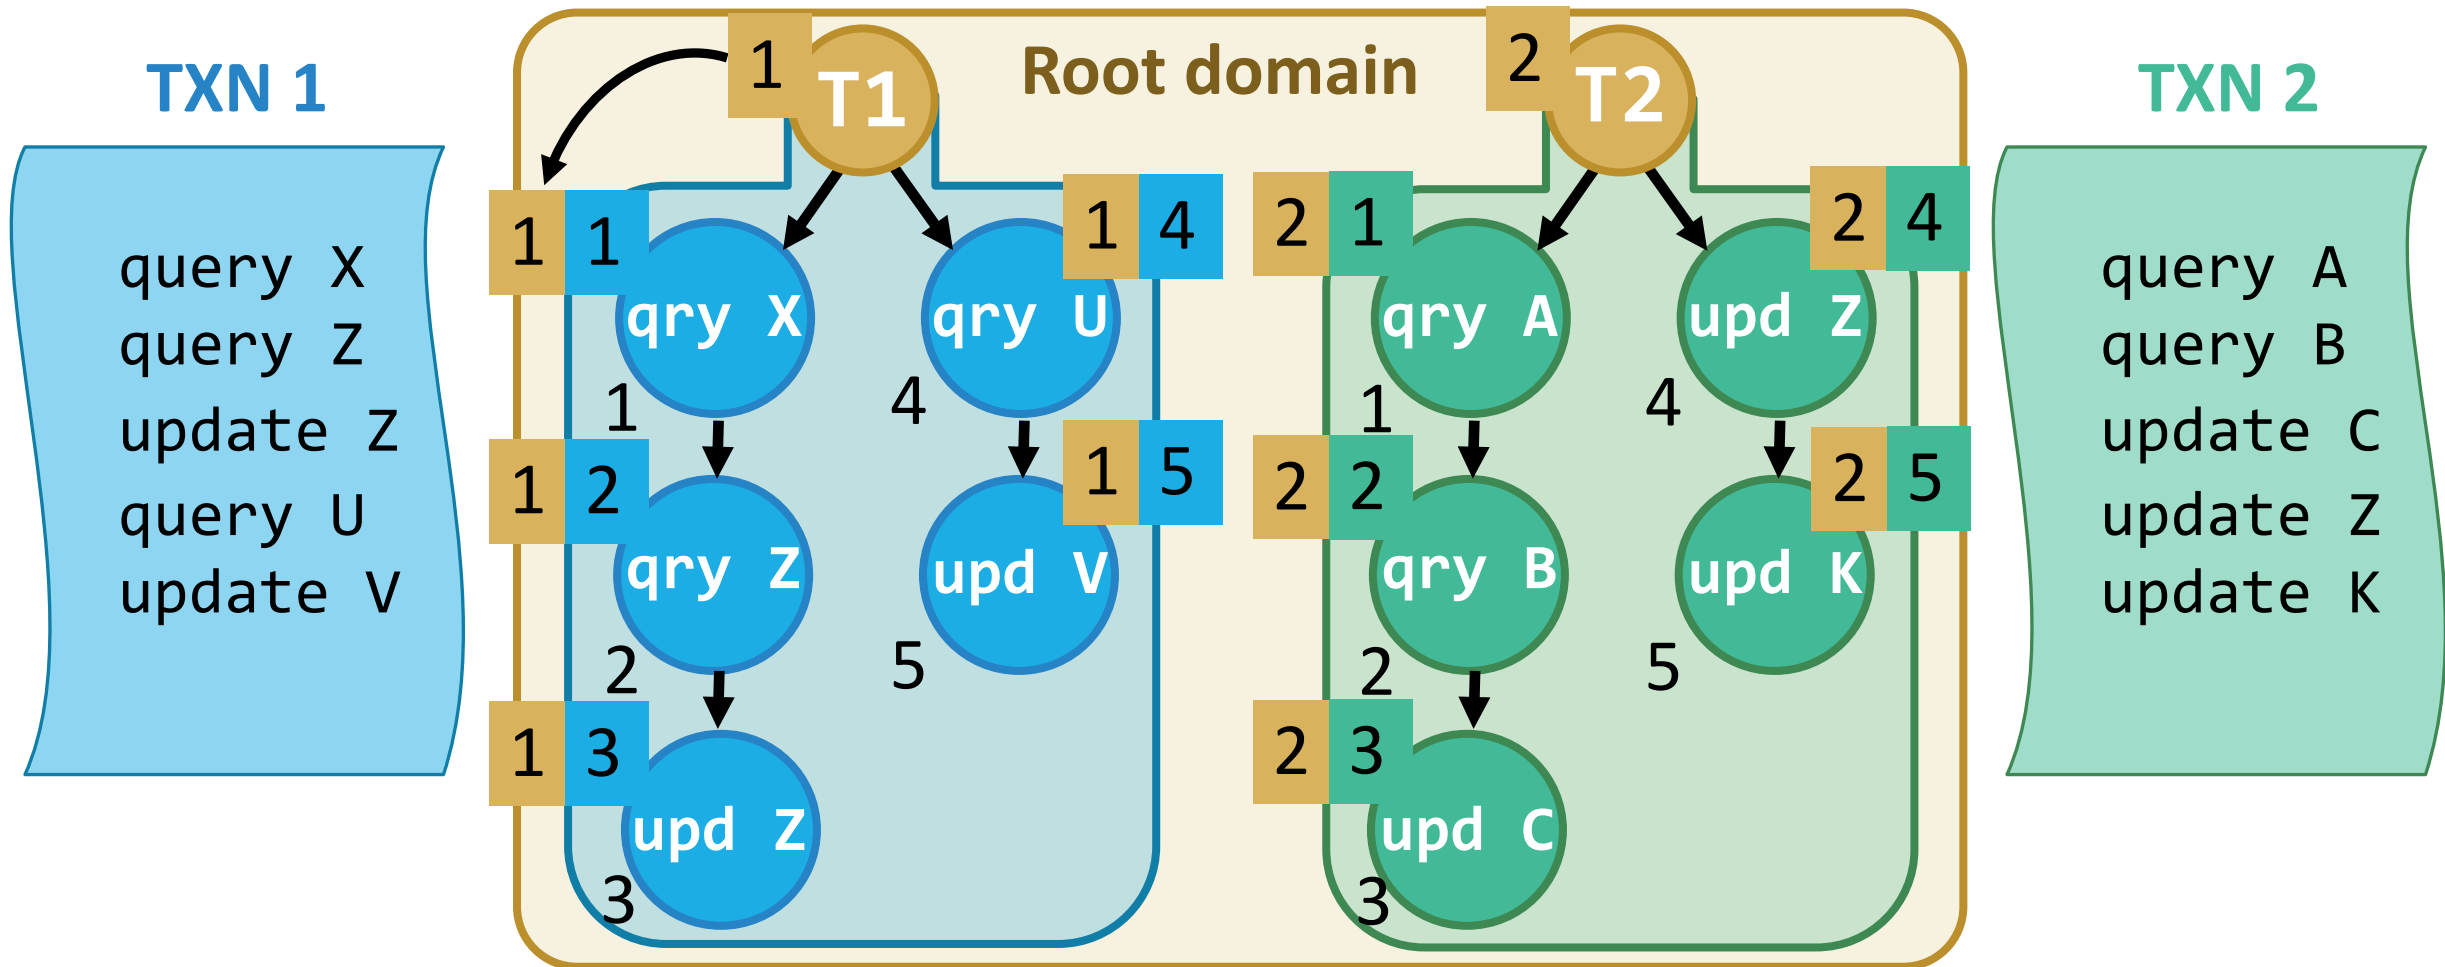
---

**Fractal** assigns a *fractal virtual time (VT)* to each task

Captures the ordering of tasks across domains, within a domain

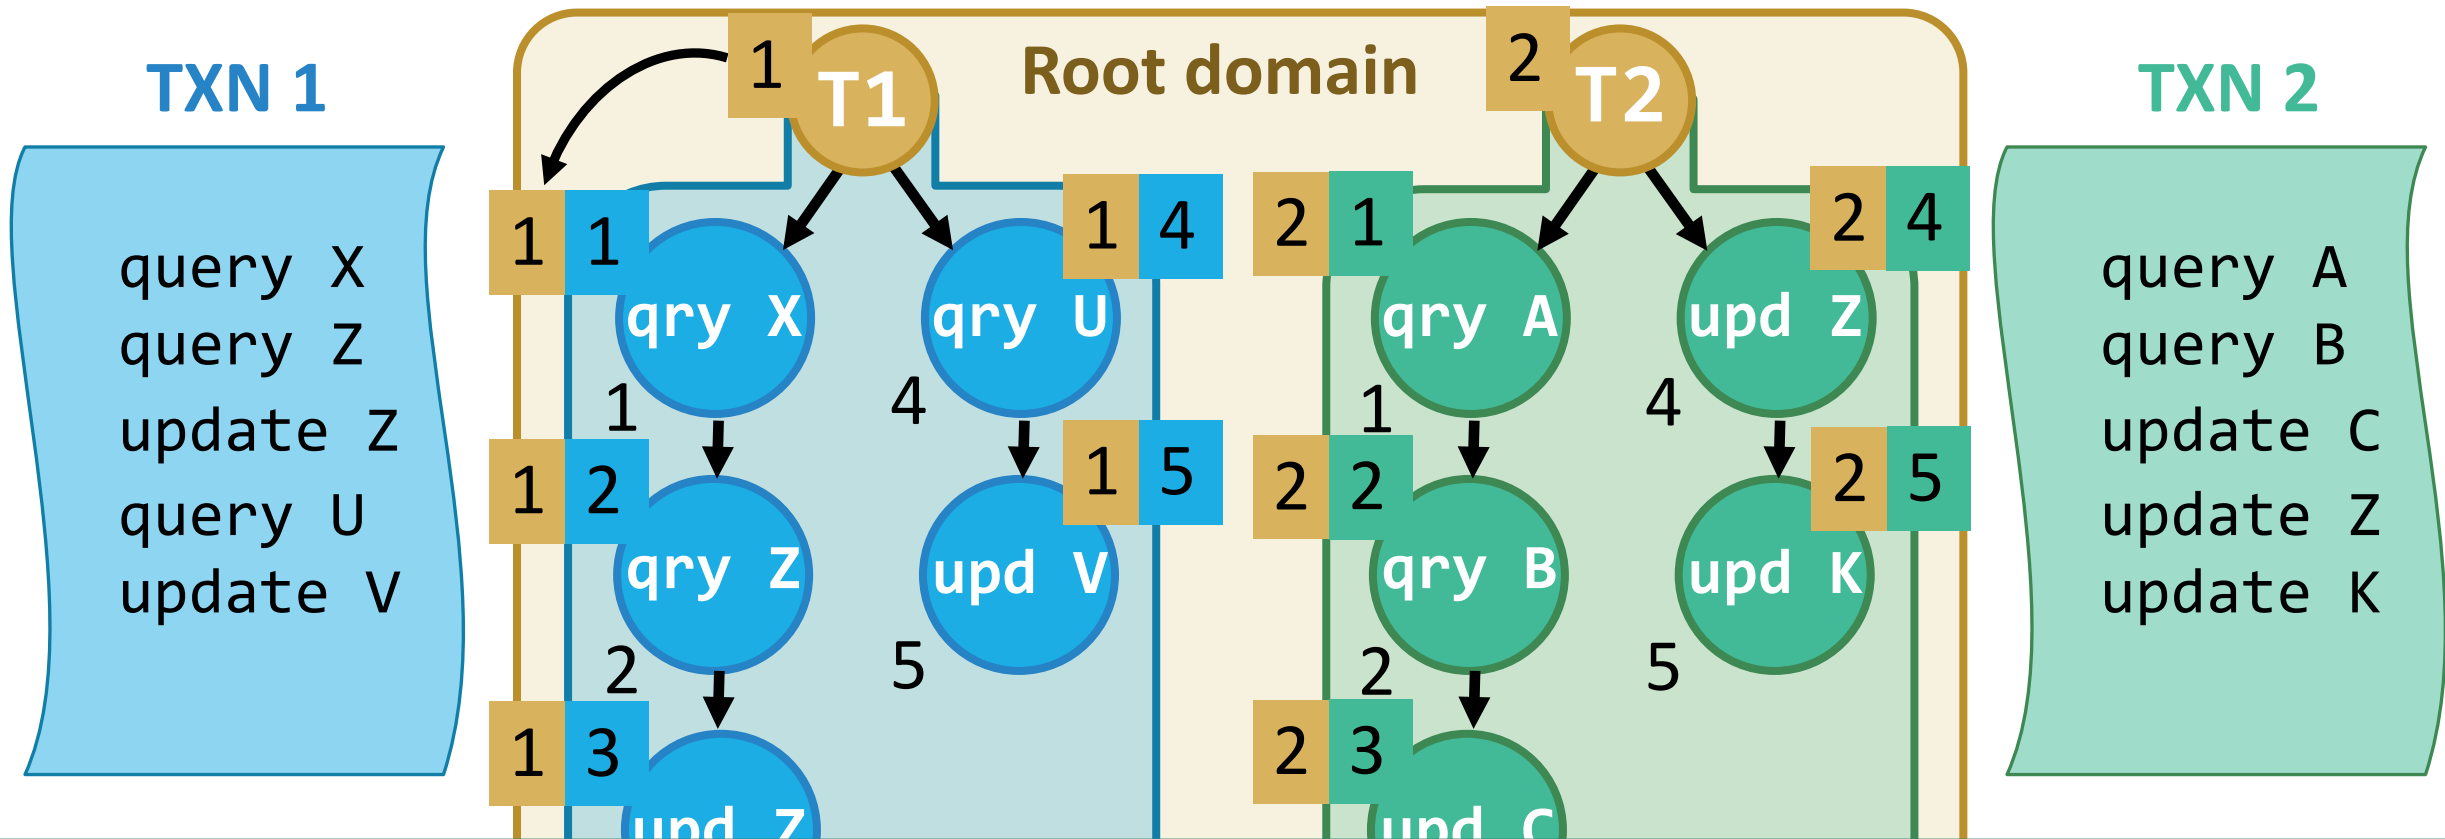


# Example: Database transactions in Fractal





# Example: Database transactions in Fractal



Fractal VT captures all ordering information

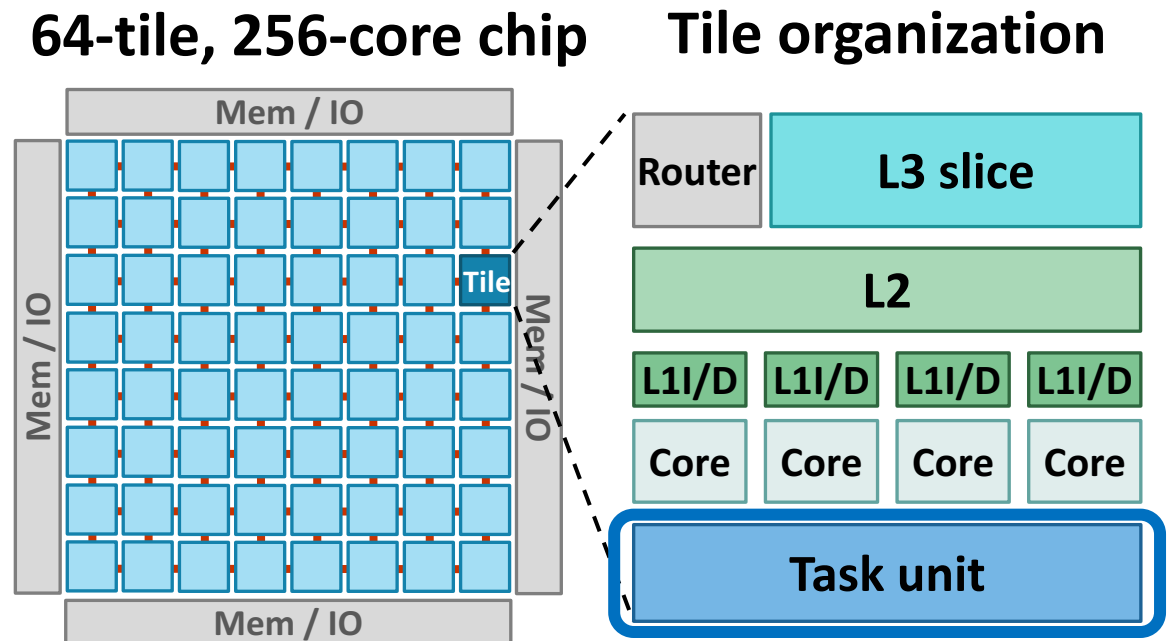
# Swarm<sub>[MICRO'15]</sub> : An efficient substrate for ordered speculation

Swarm executes tasks speculatively and out of order

Large hardware task queues

Scalable ordered commits

Scalable ordered speculation



**Efficiently supports tiny speculative tasks**

# Fractal features

---


**Fractal** VT construction requires no centralized structures

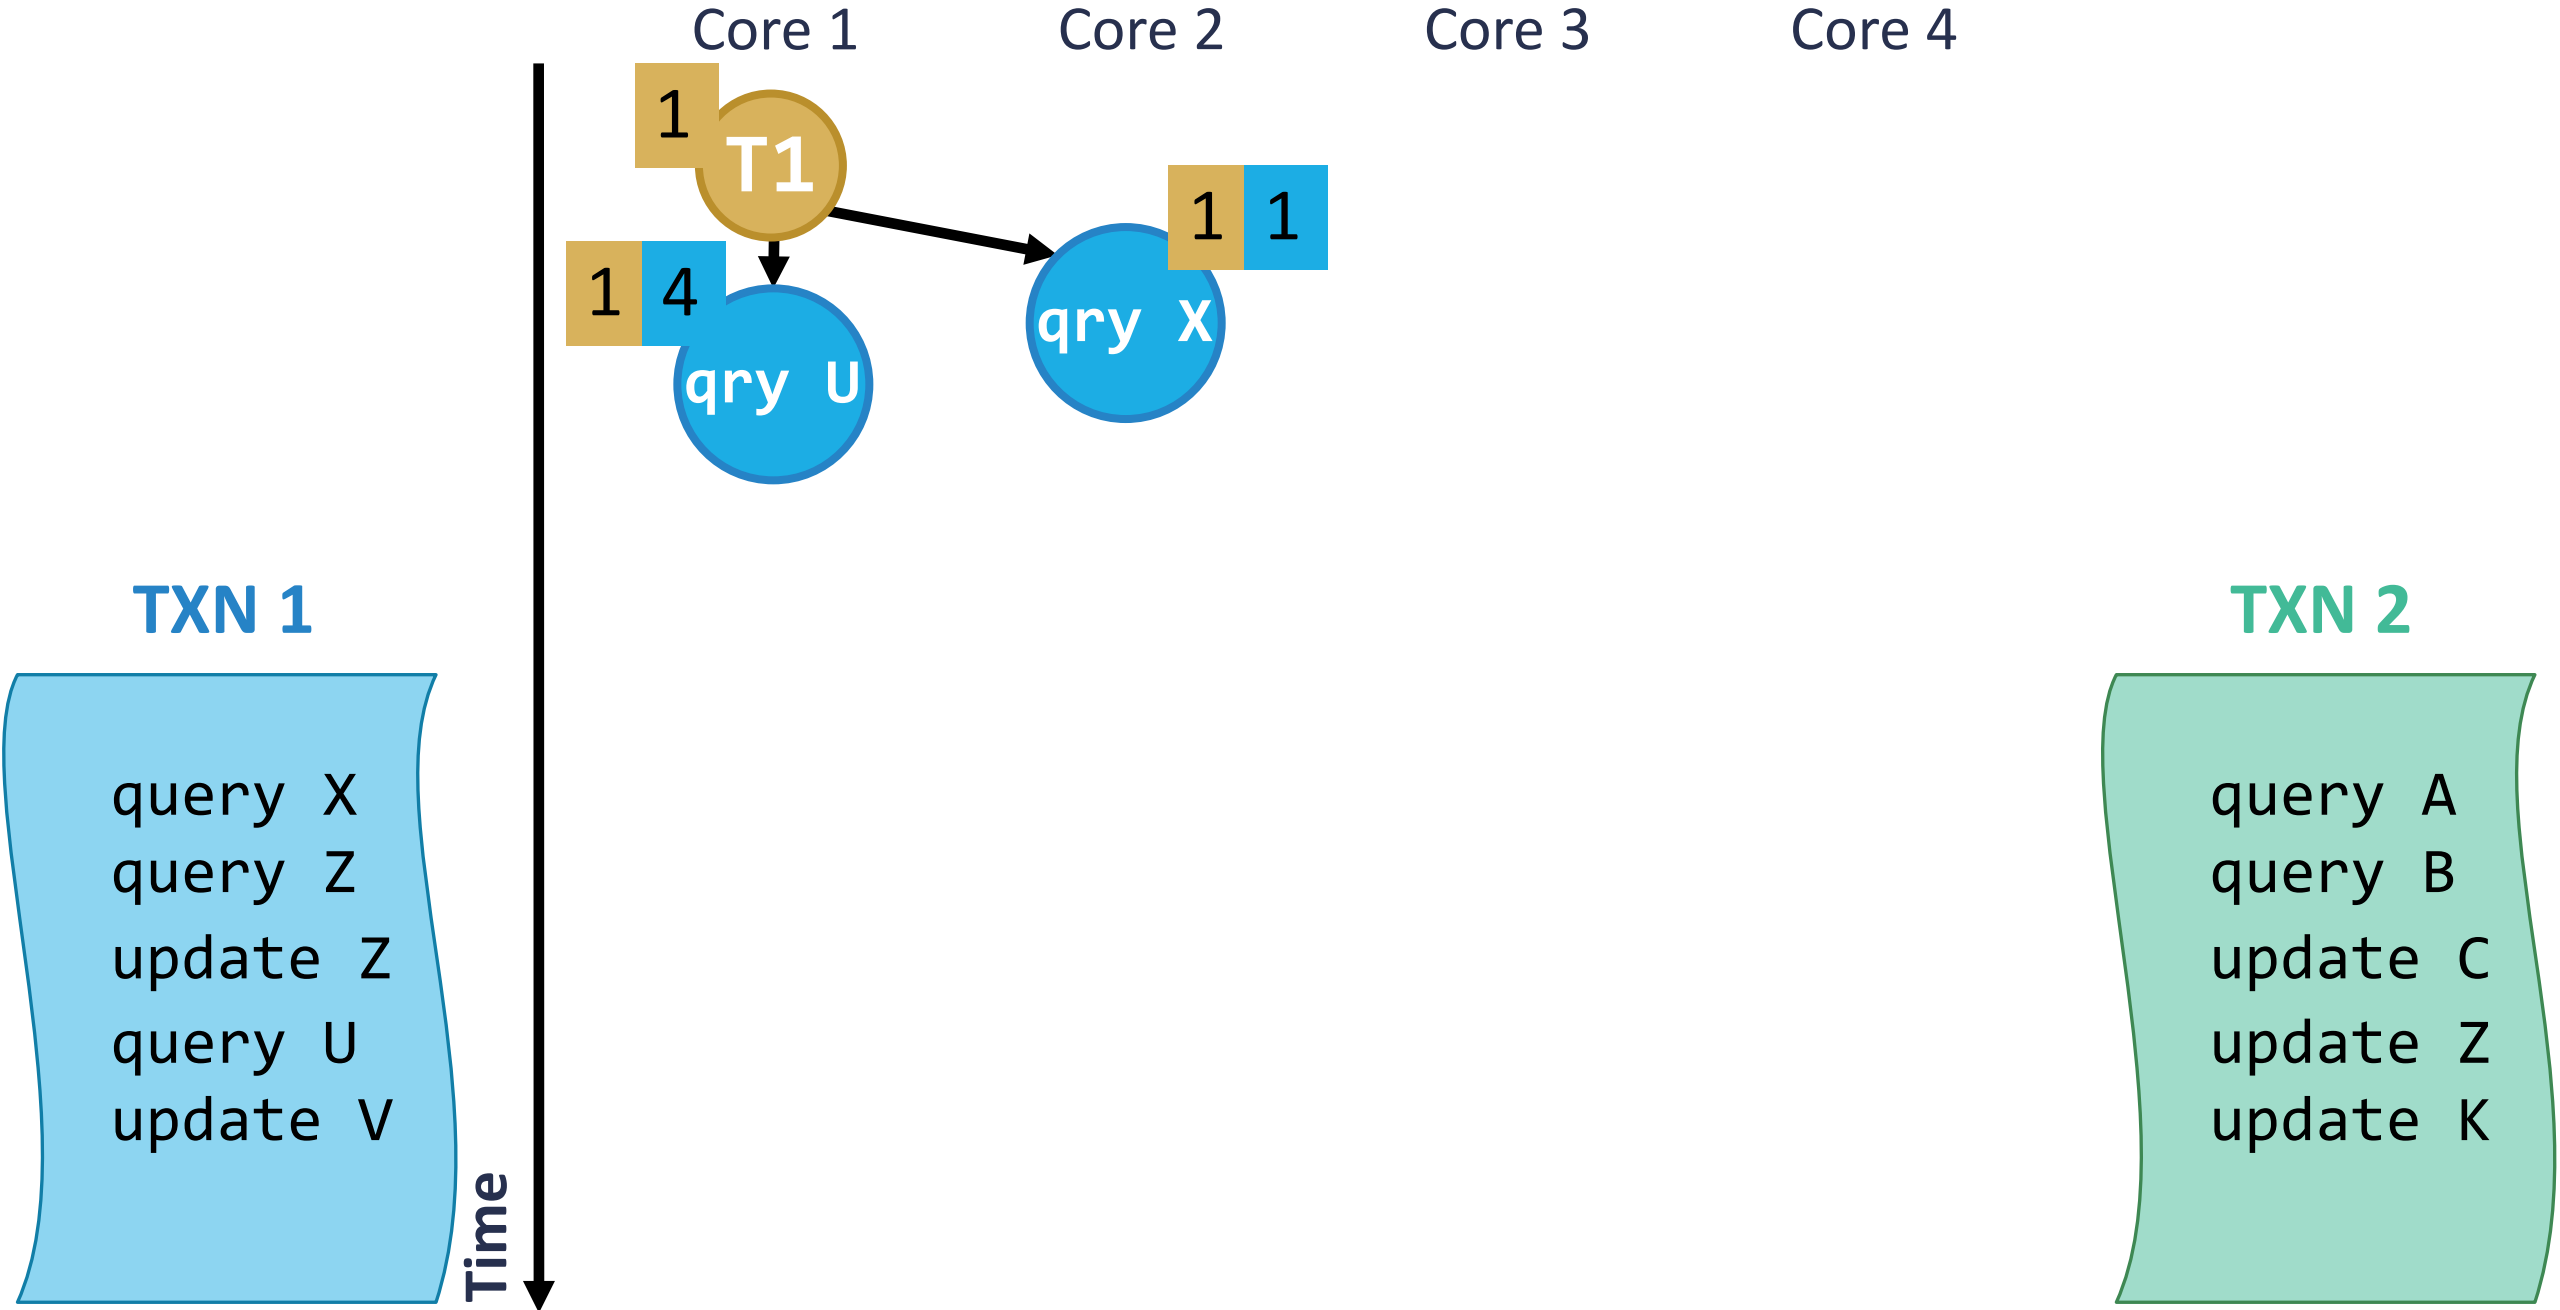
**Fractal** VT assigns order dynamically

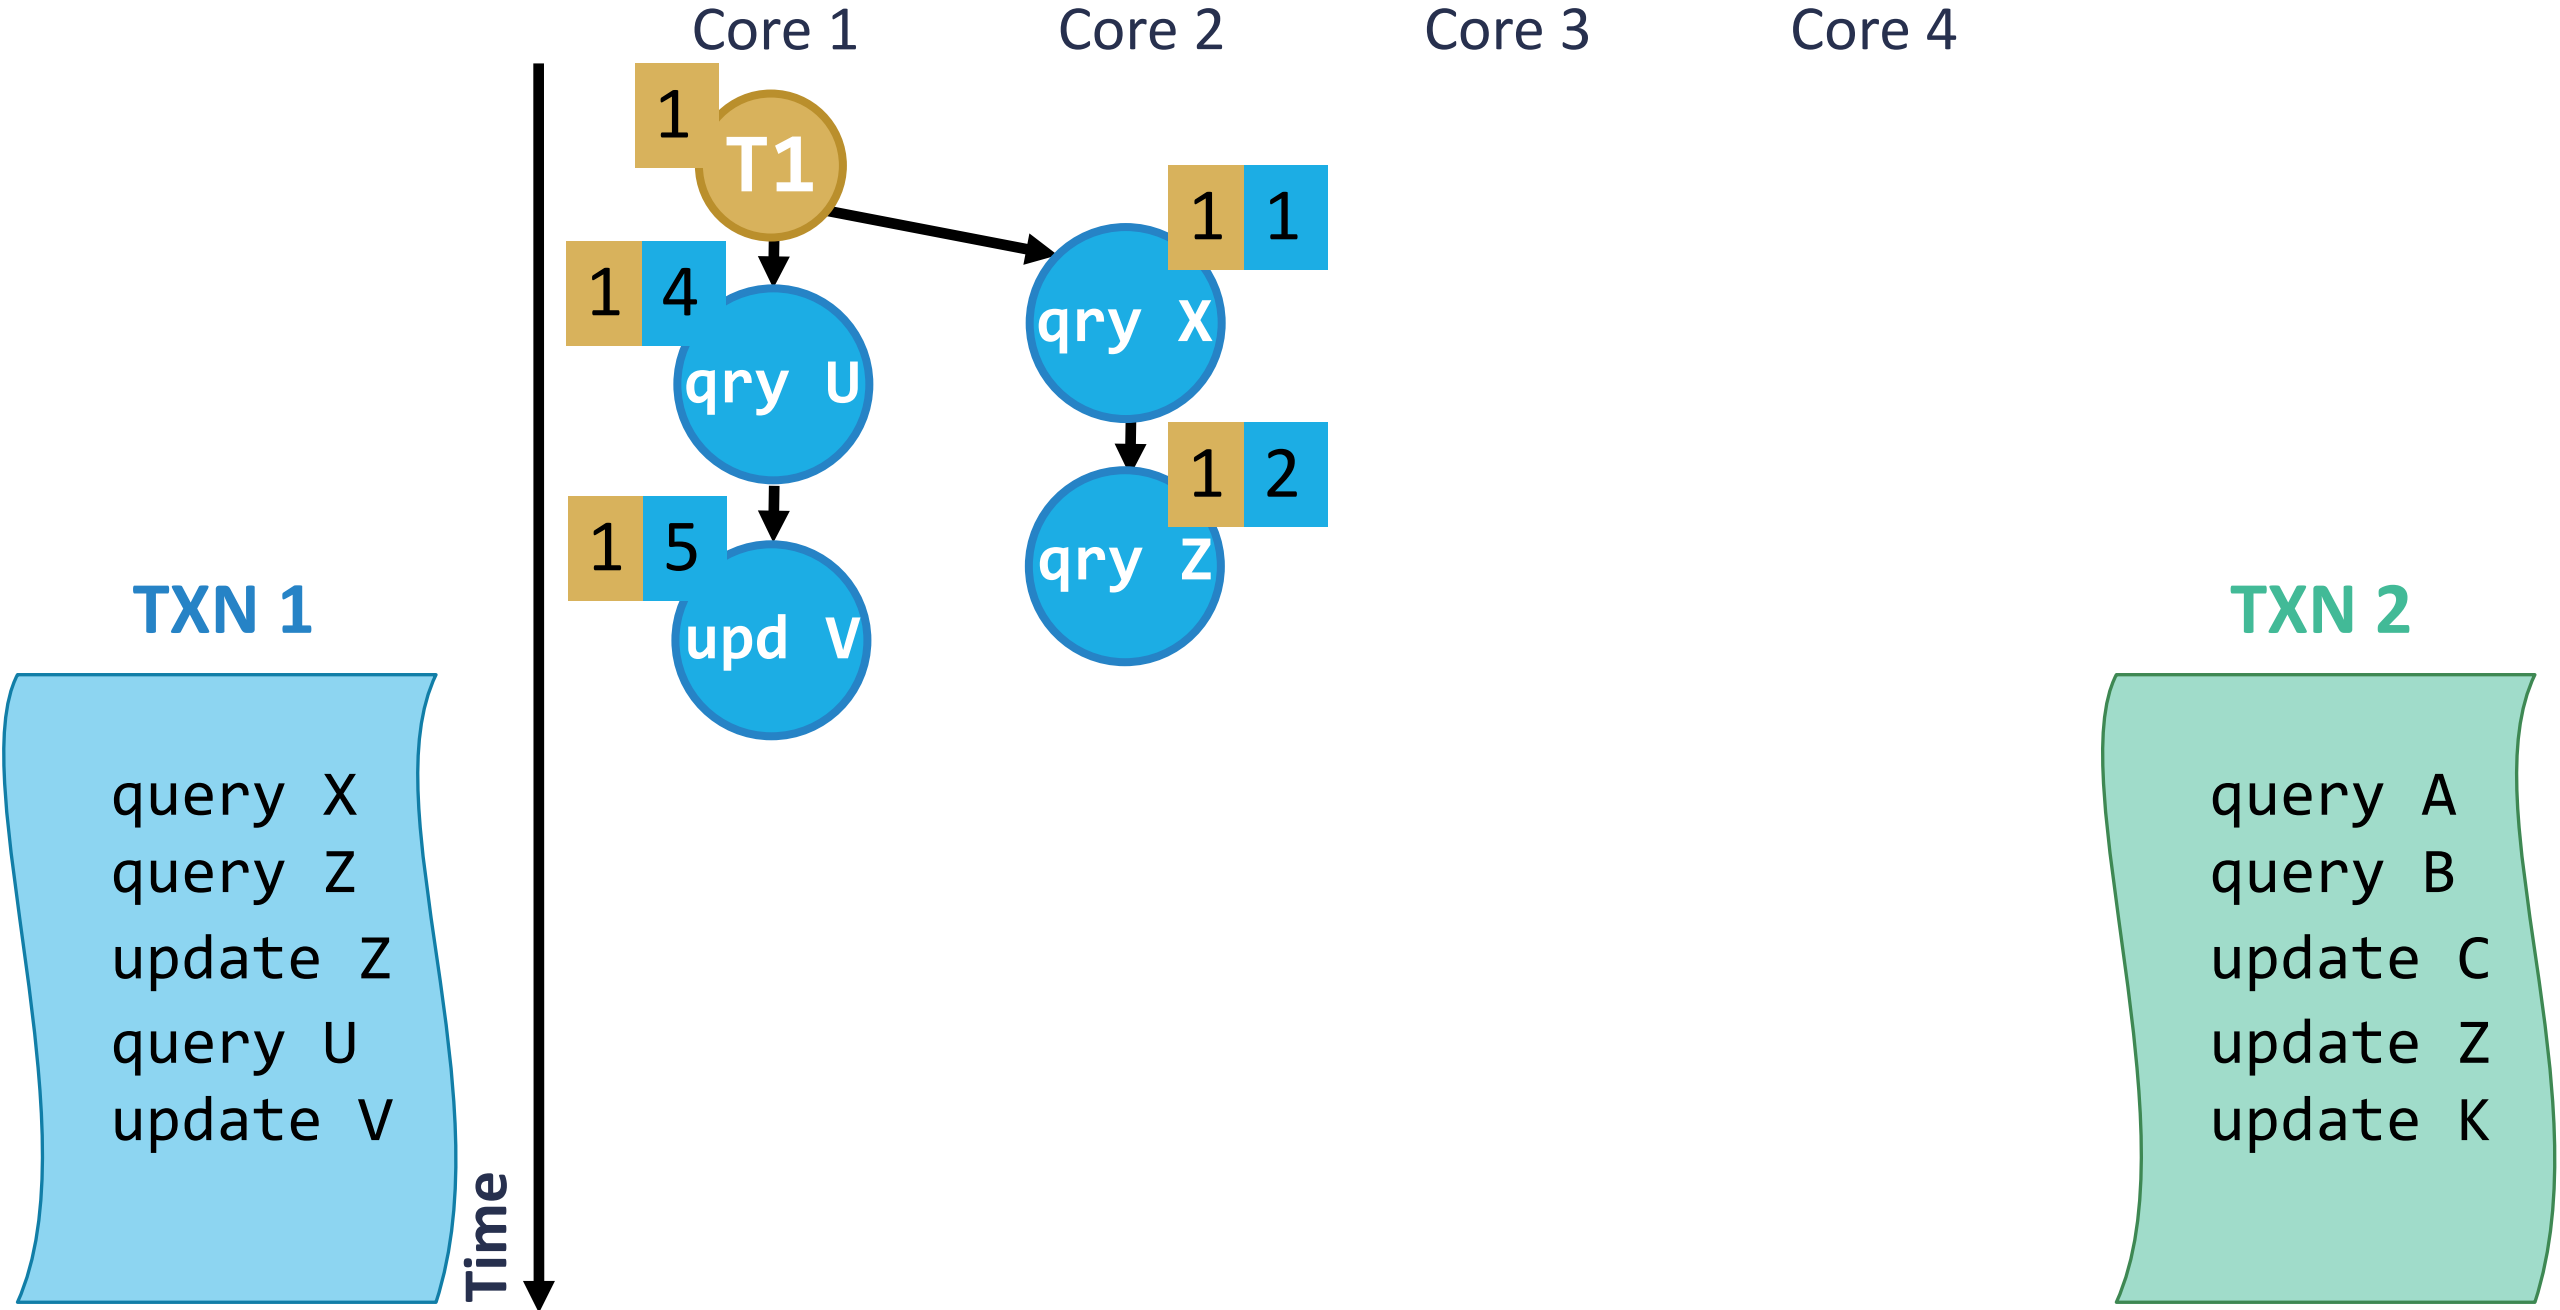
Hardware supports a few number of *concurrent depths*

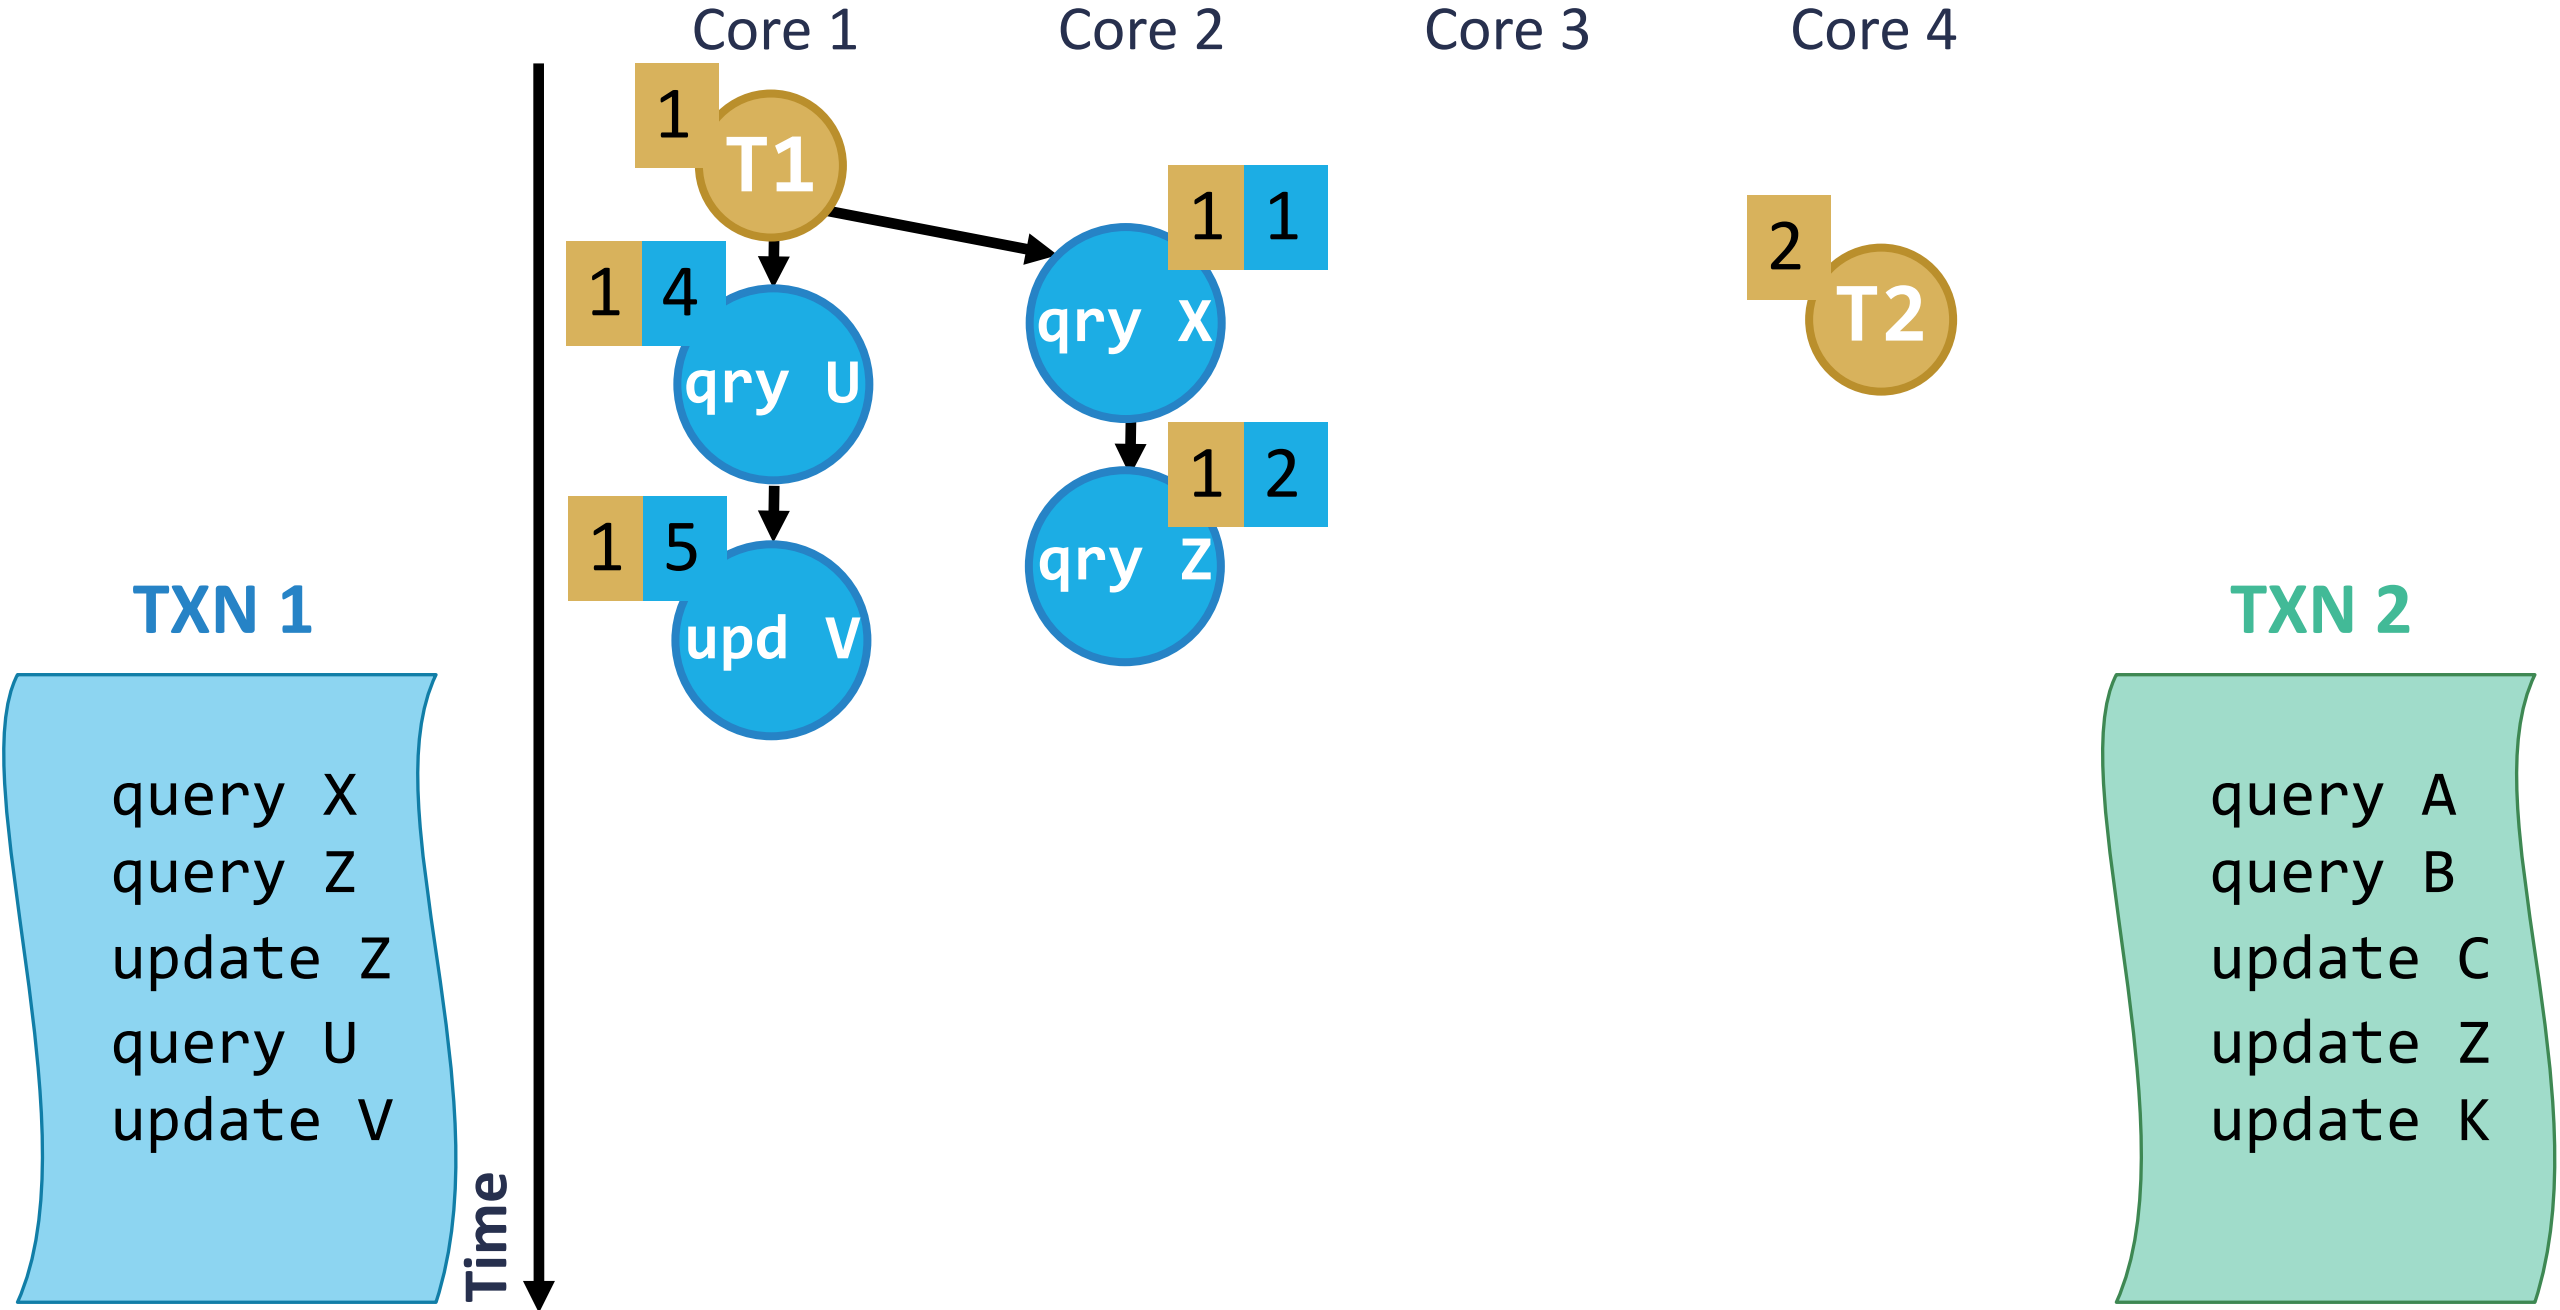
- “Zooming” operations allow for unbounded nesting
- Spill tasks from shallower domains to memory
- Parallelism compounds quickly with depth

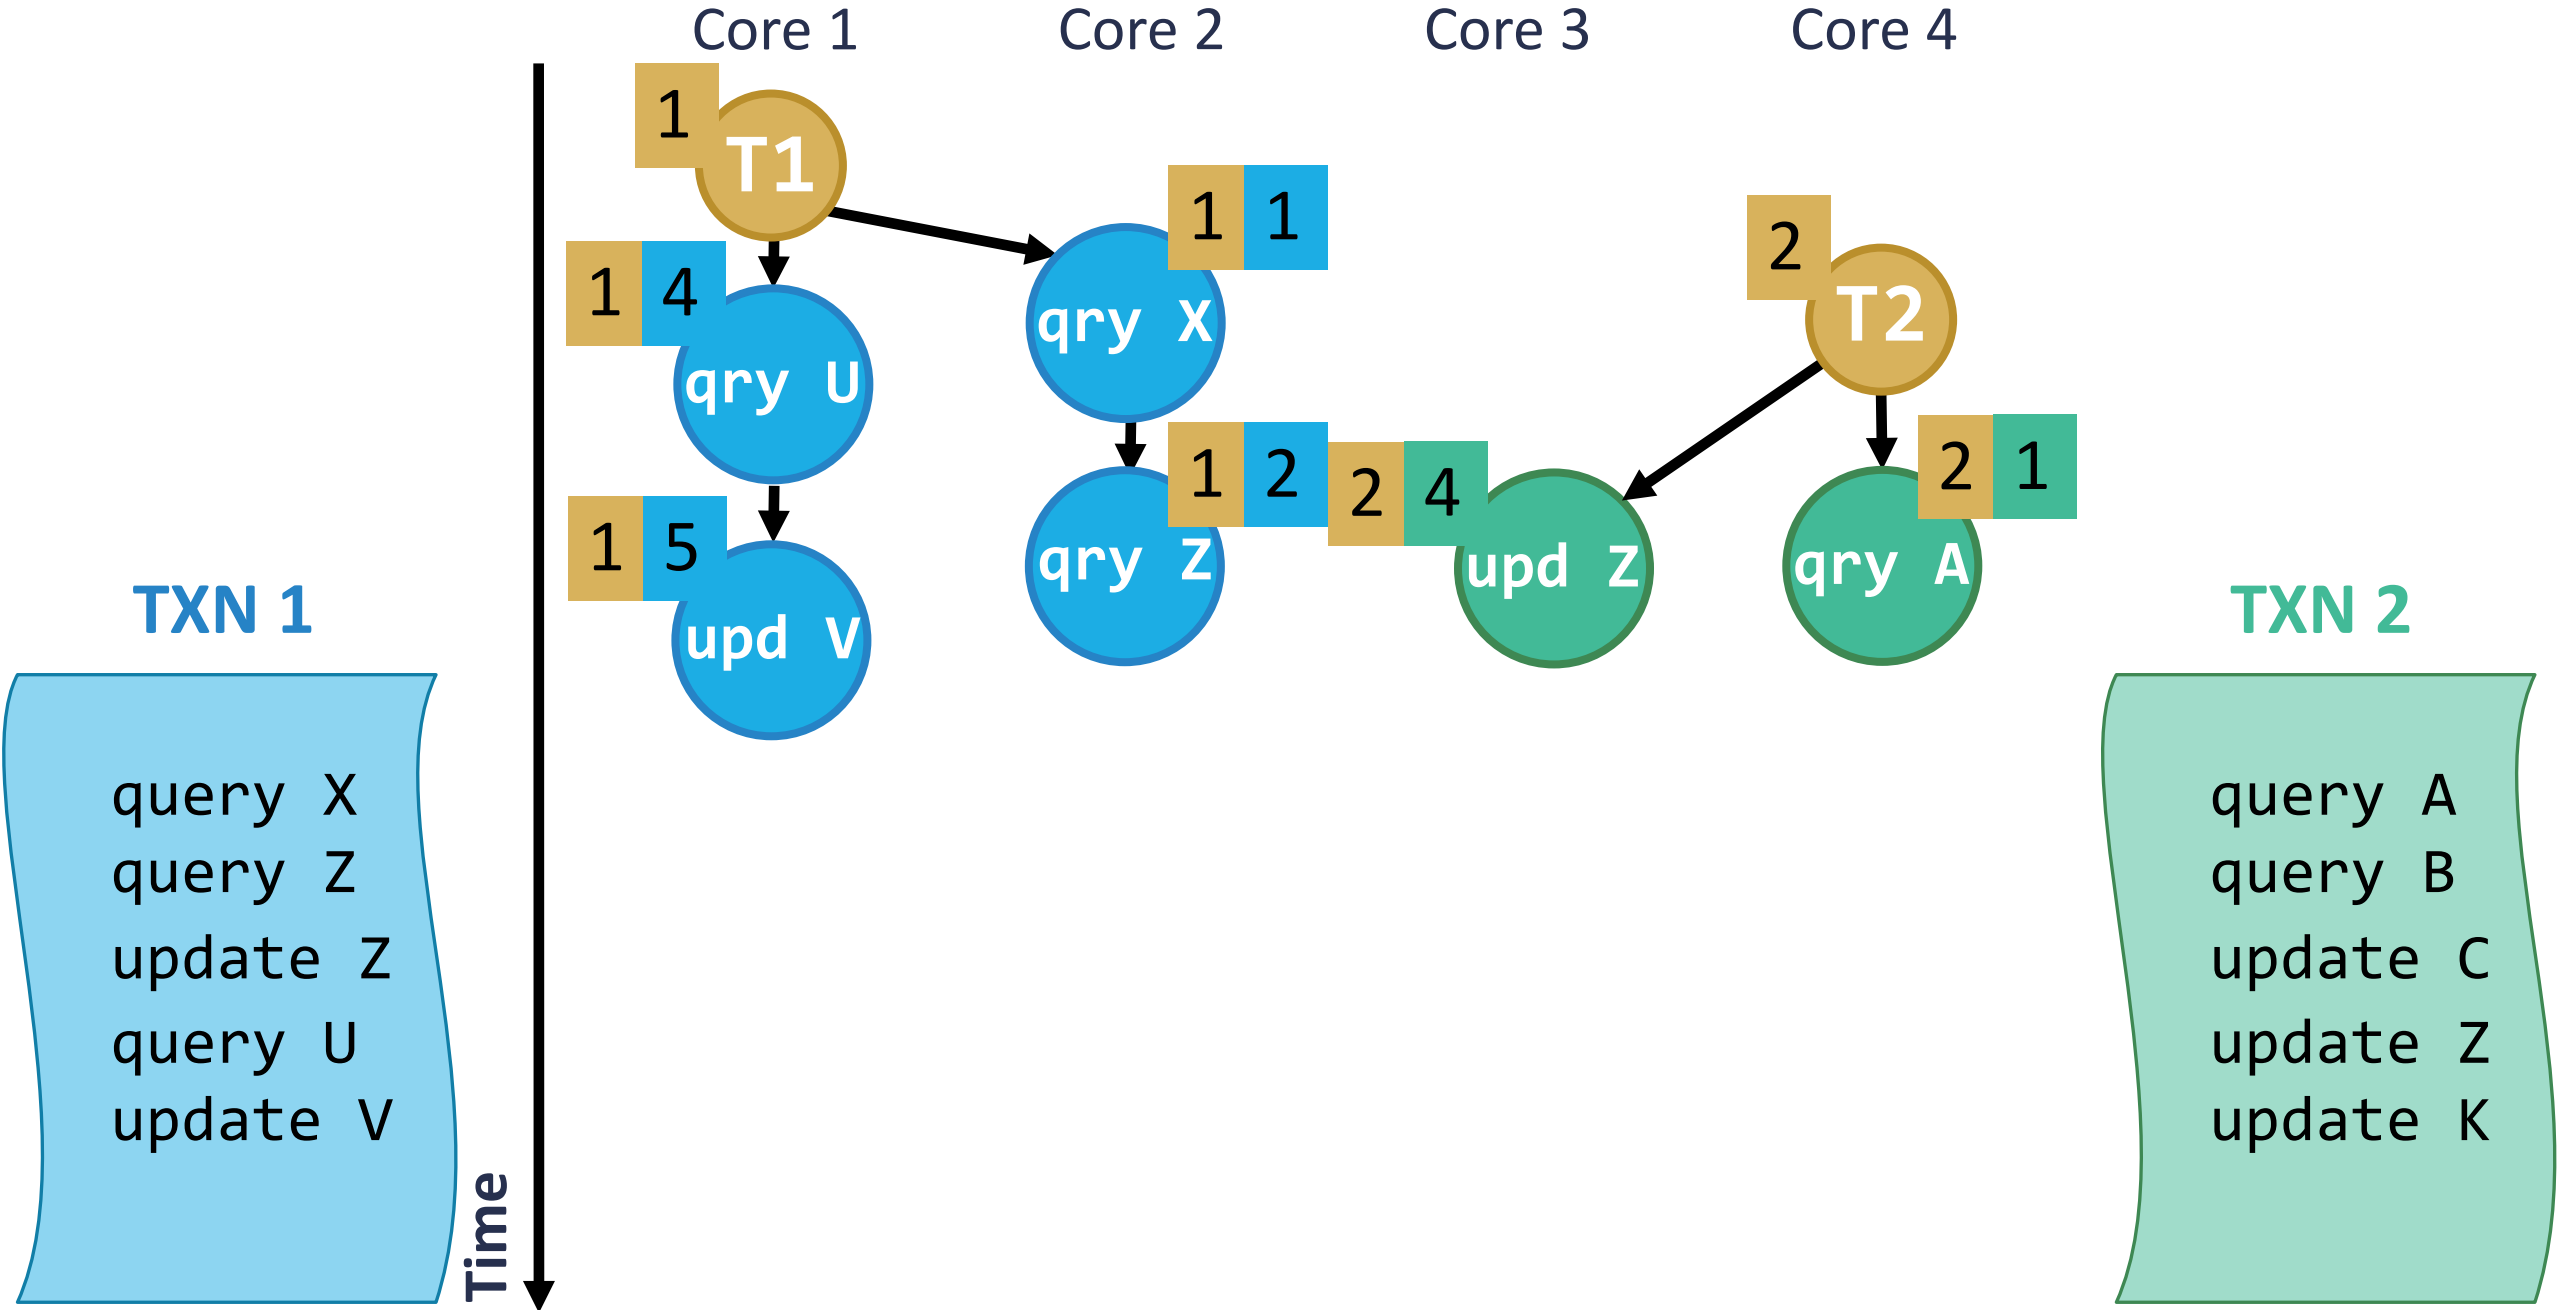
See the paper for more details!



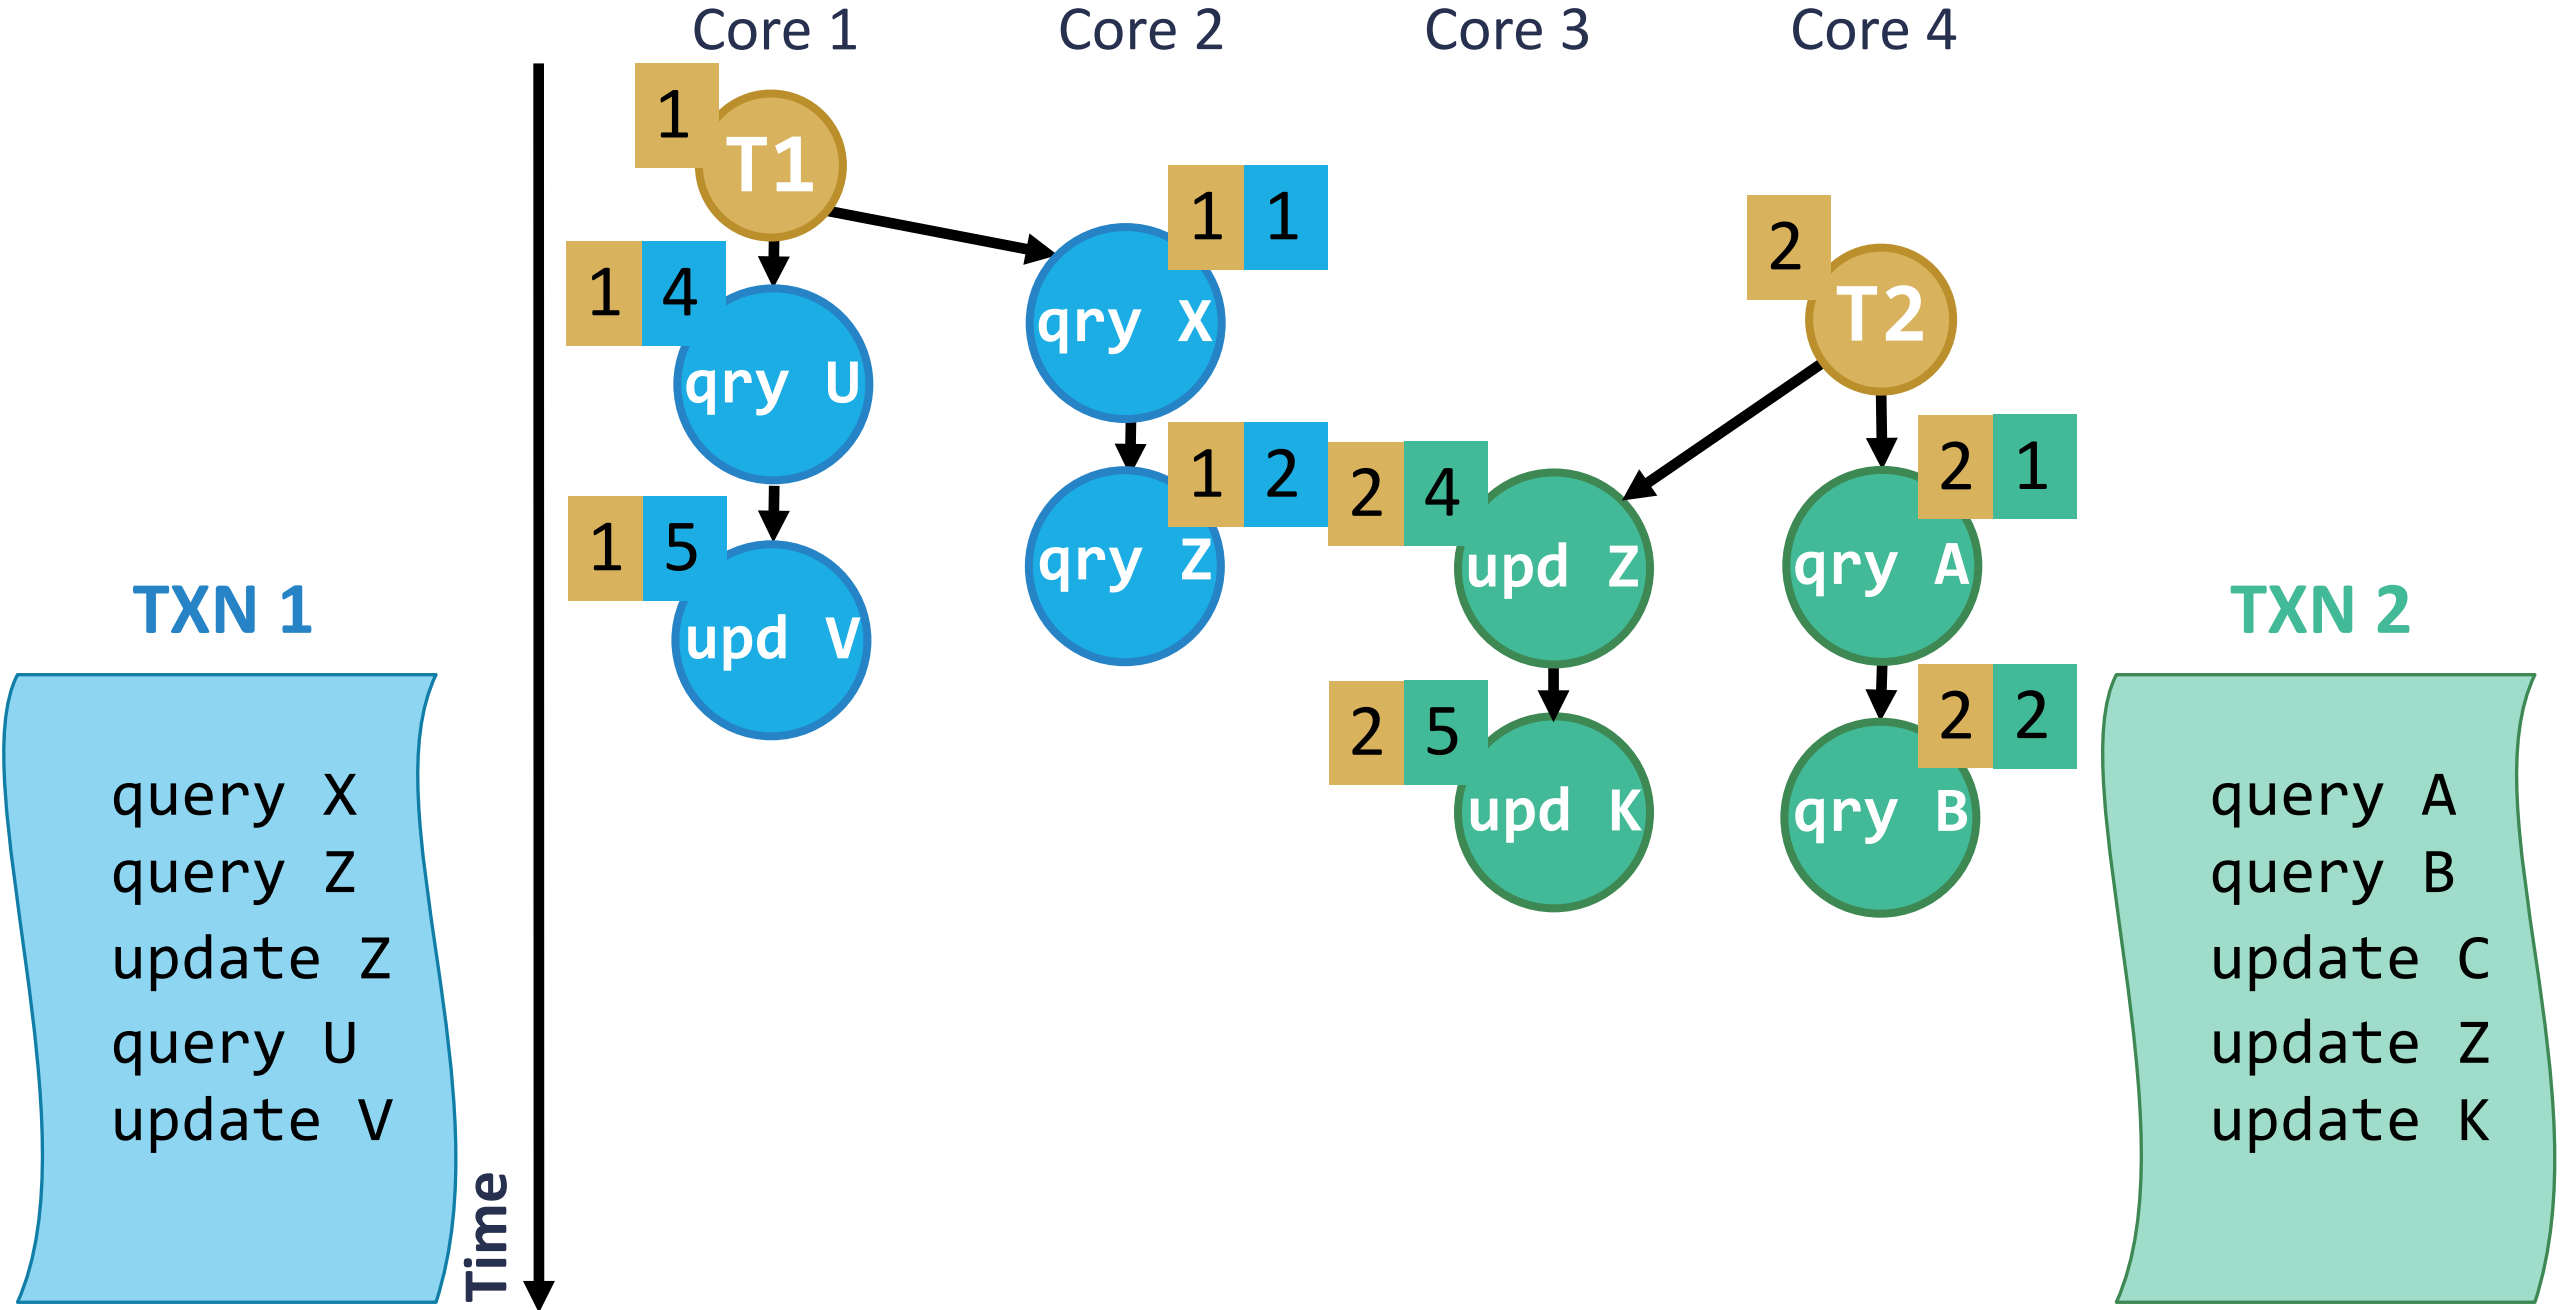


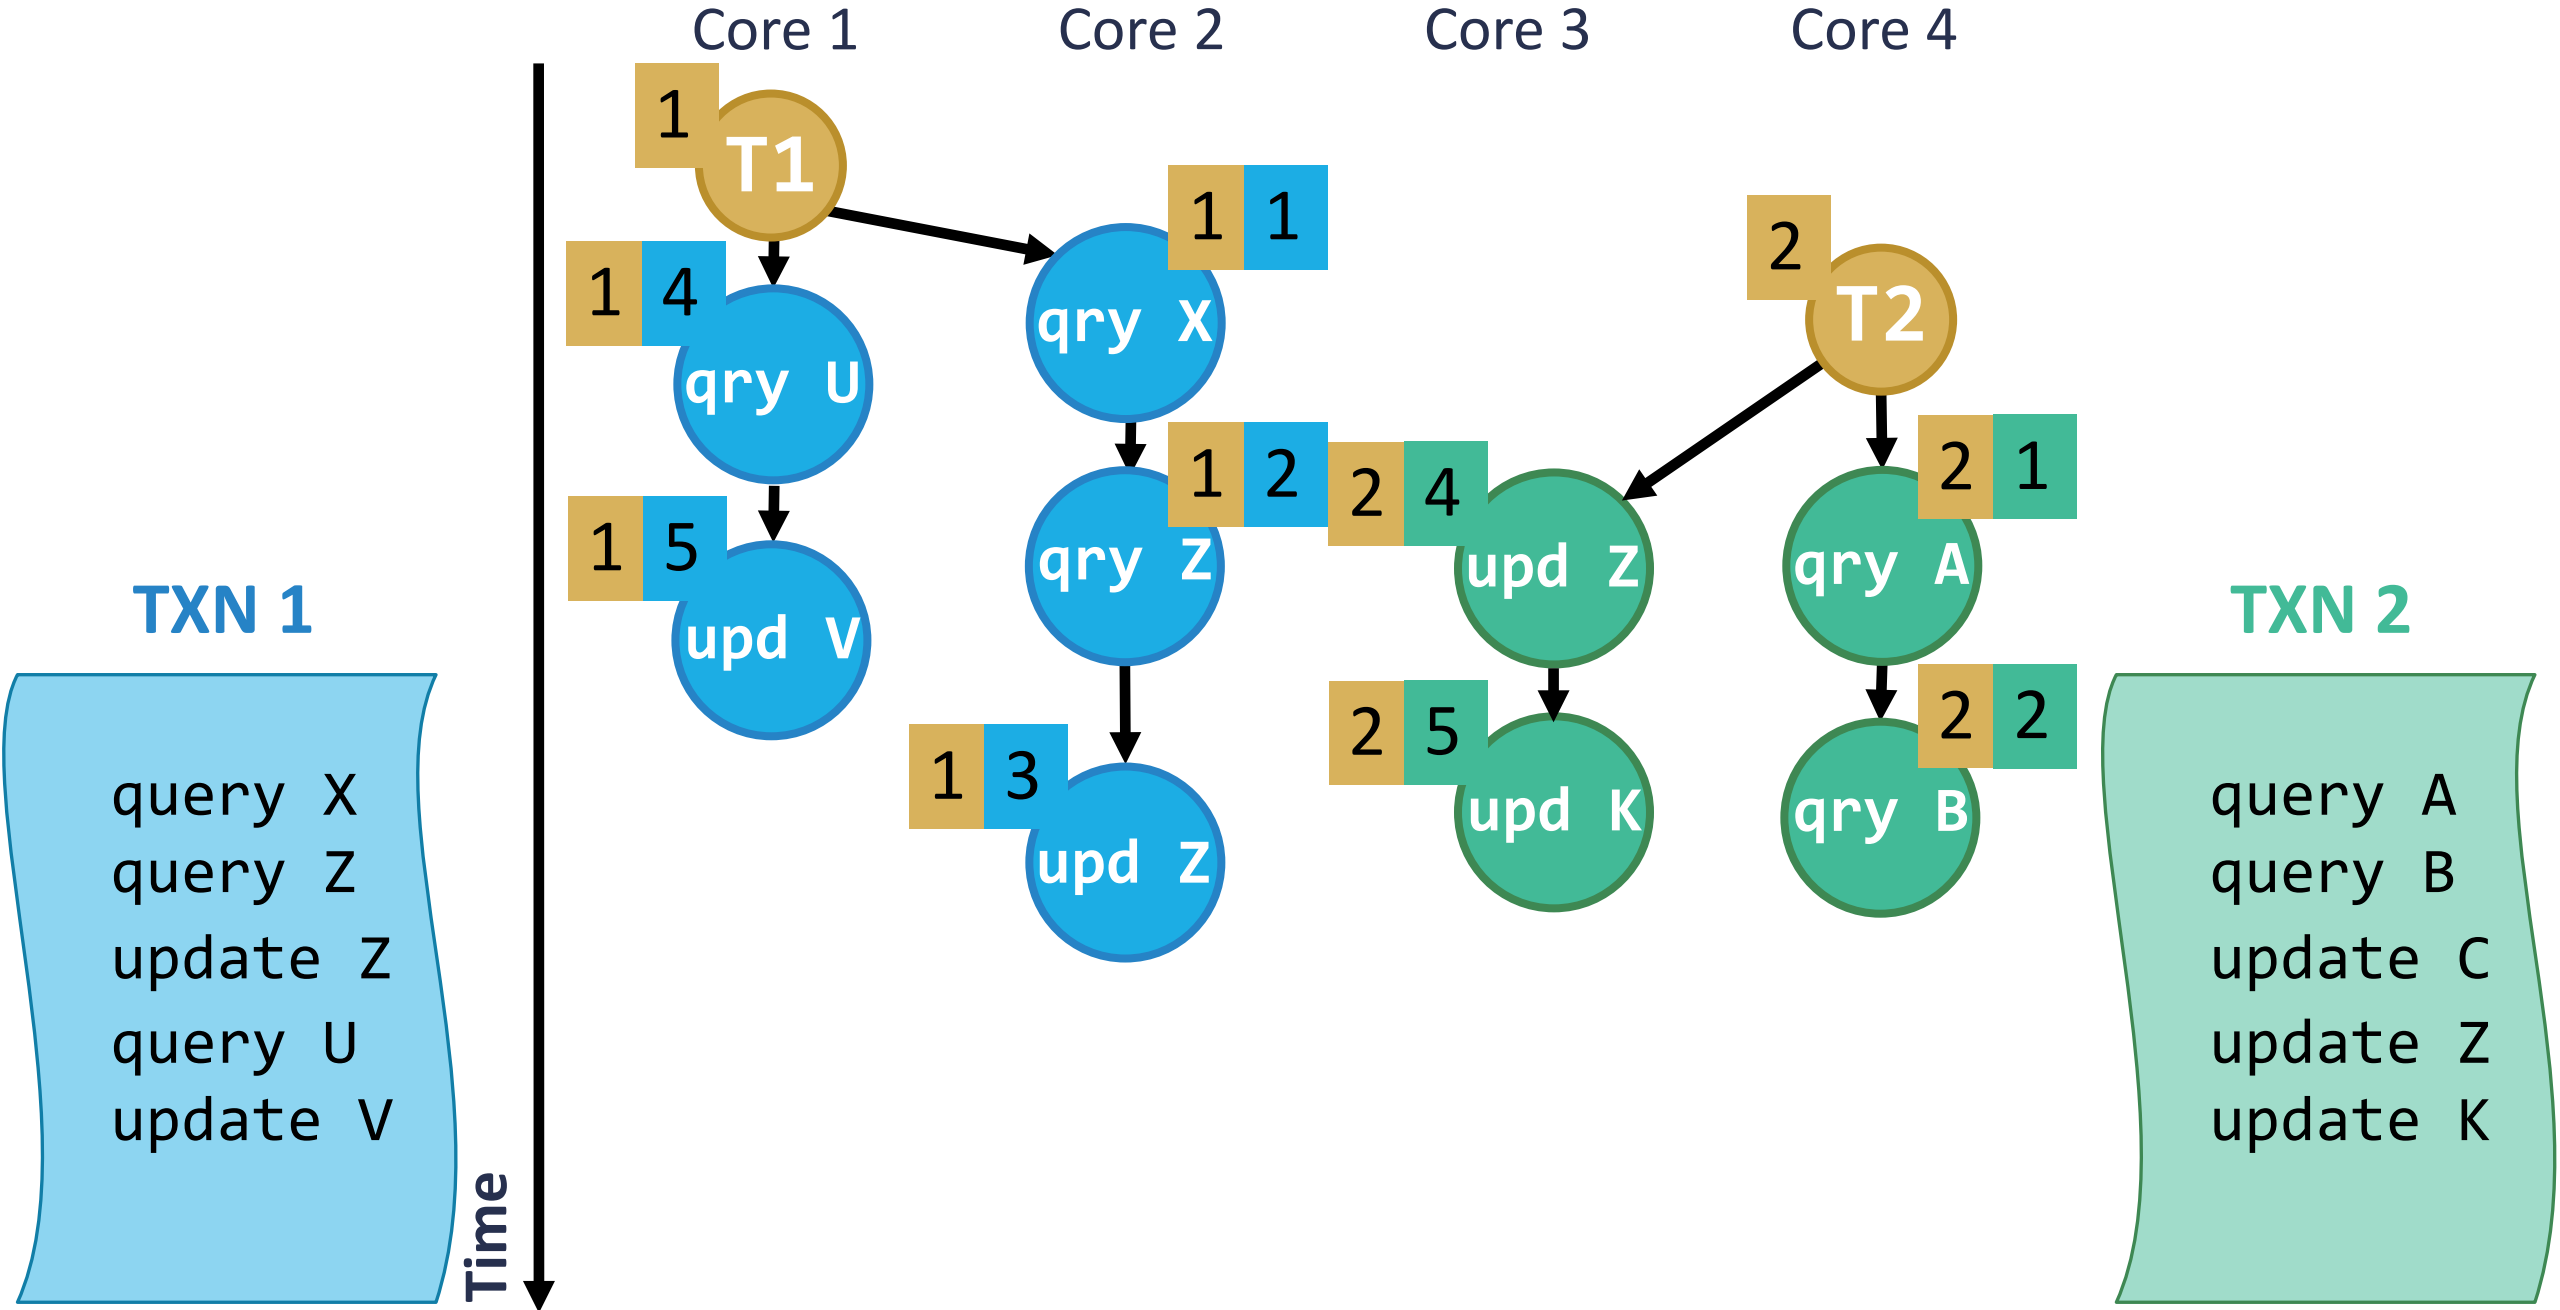


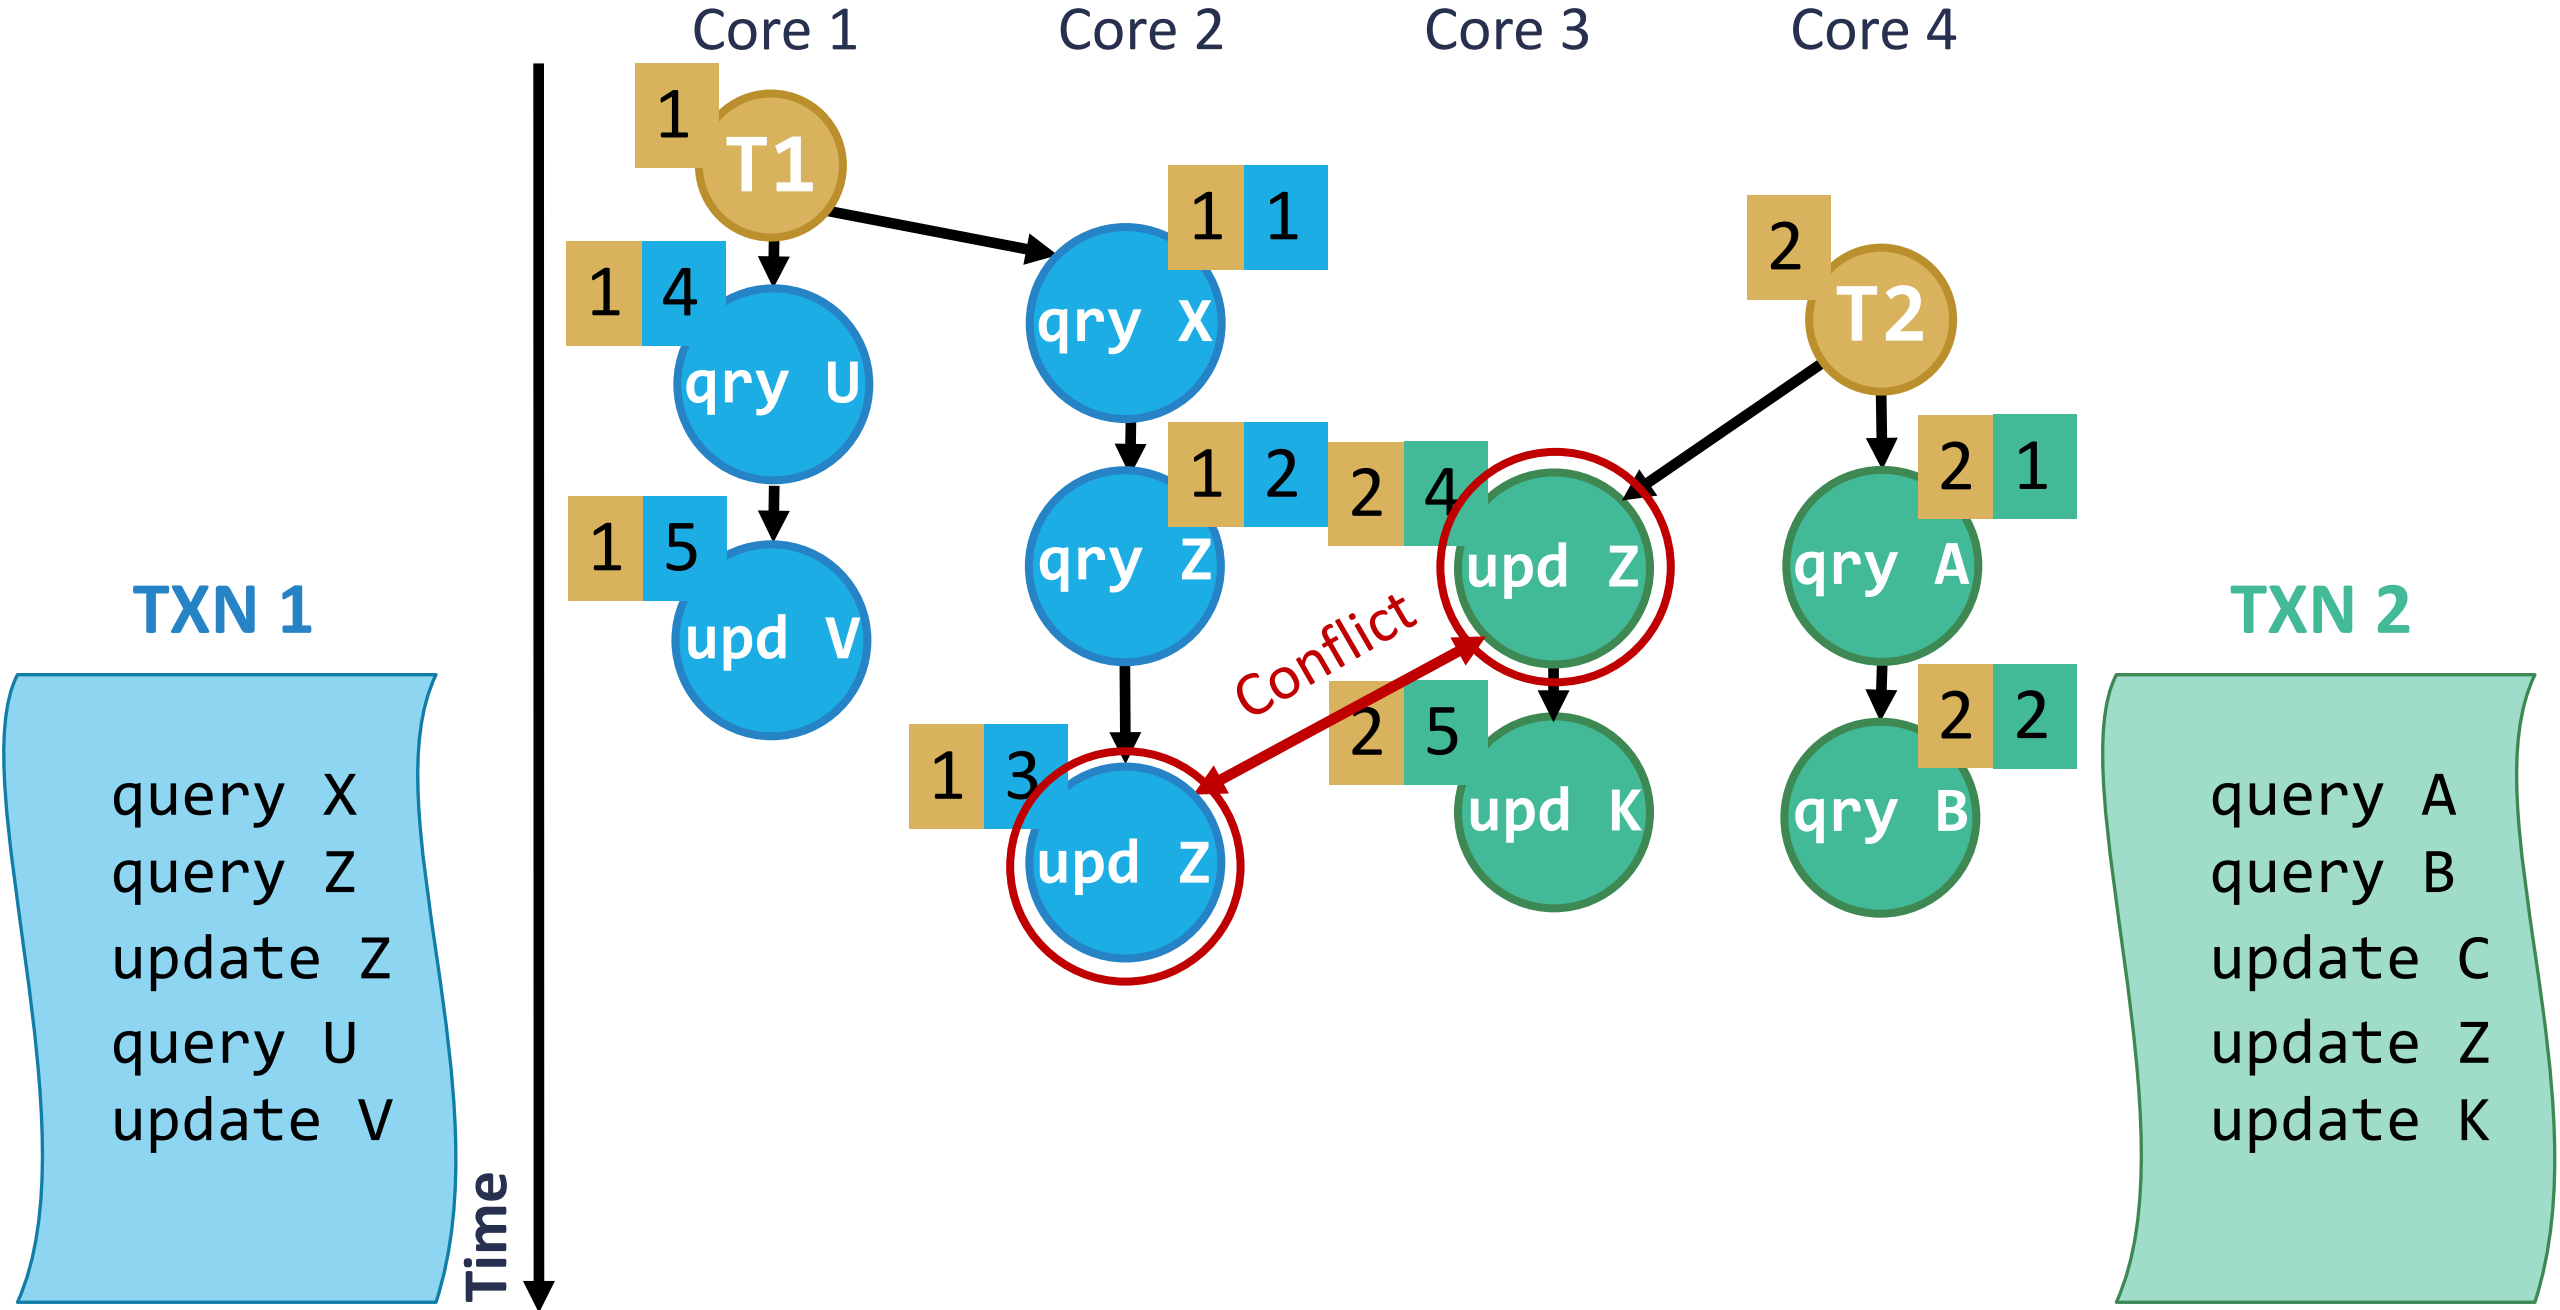


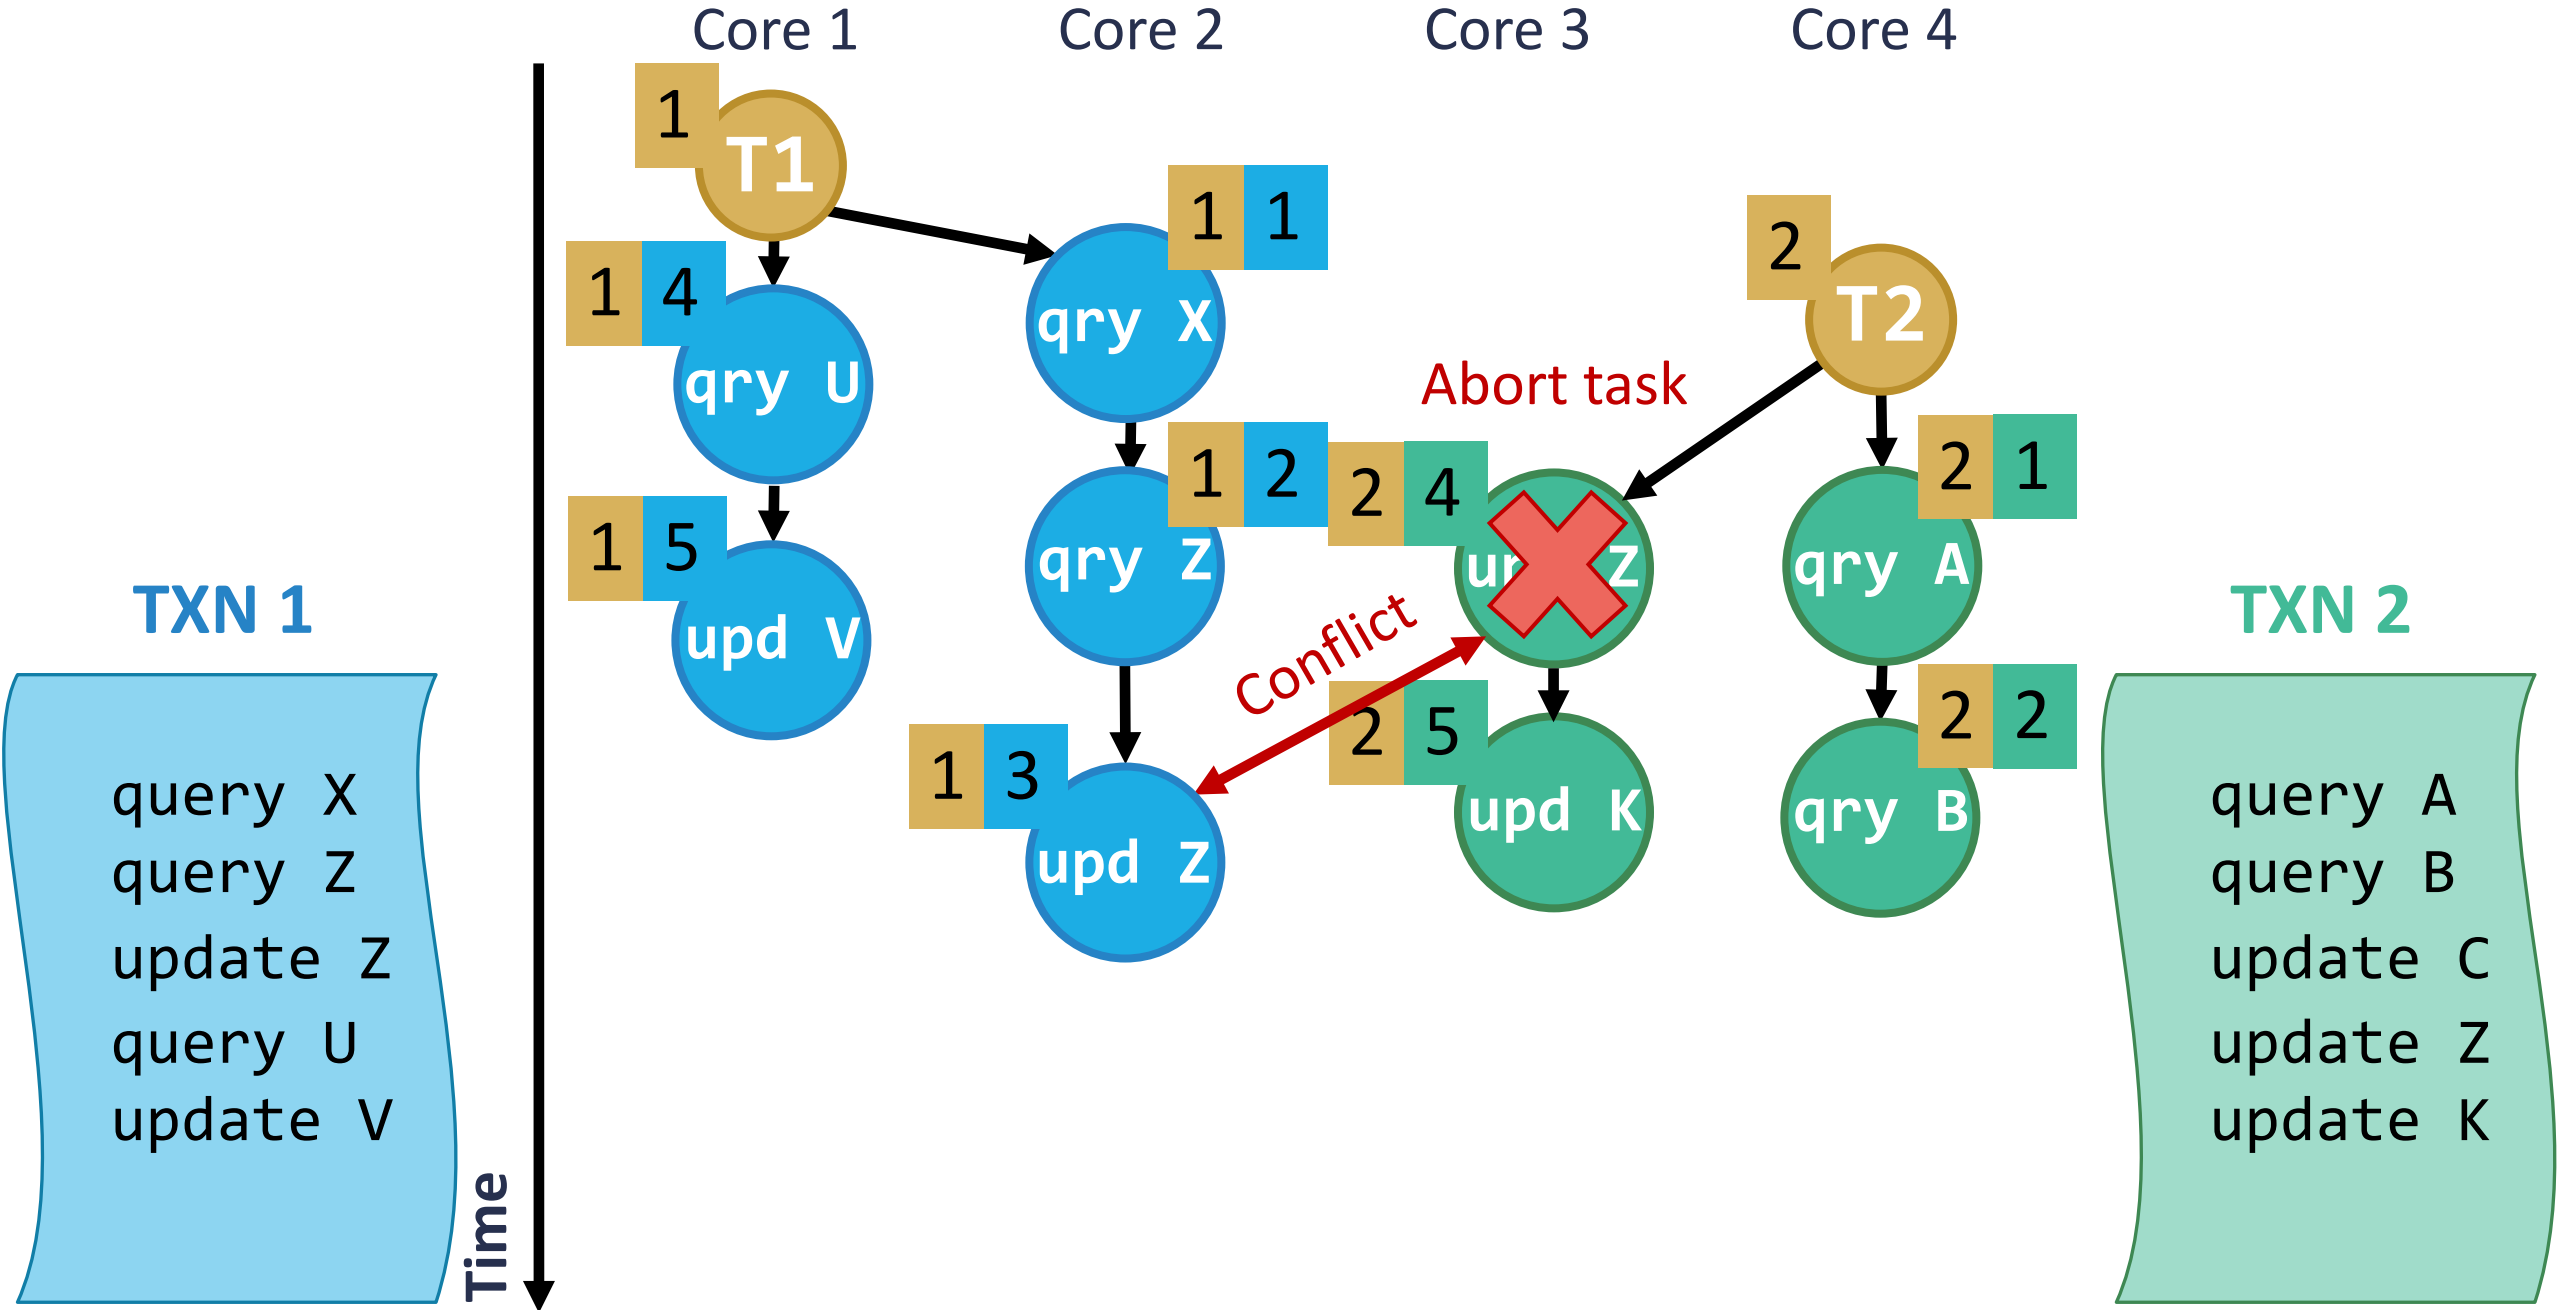


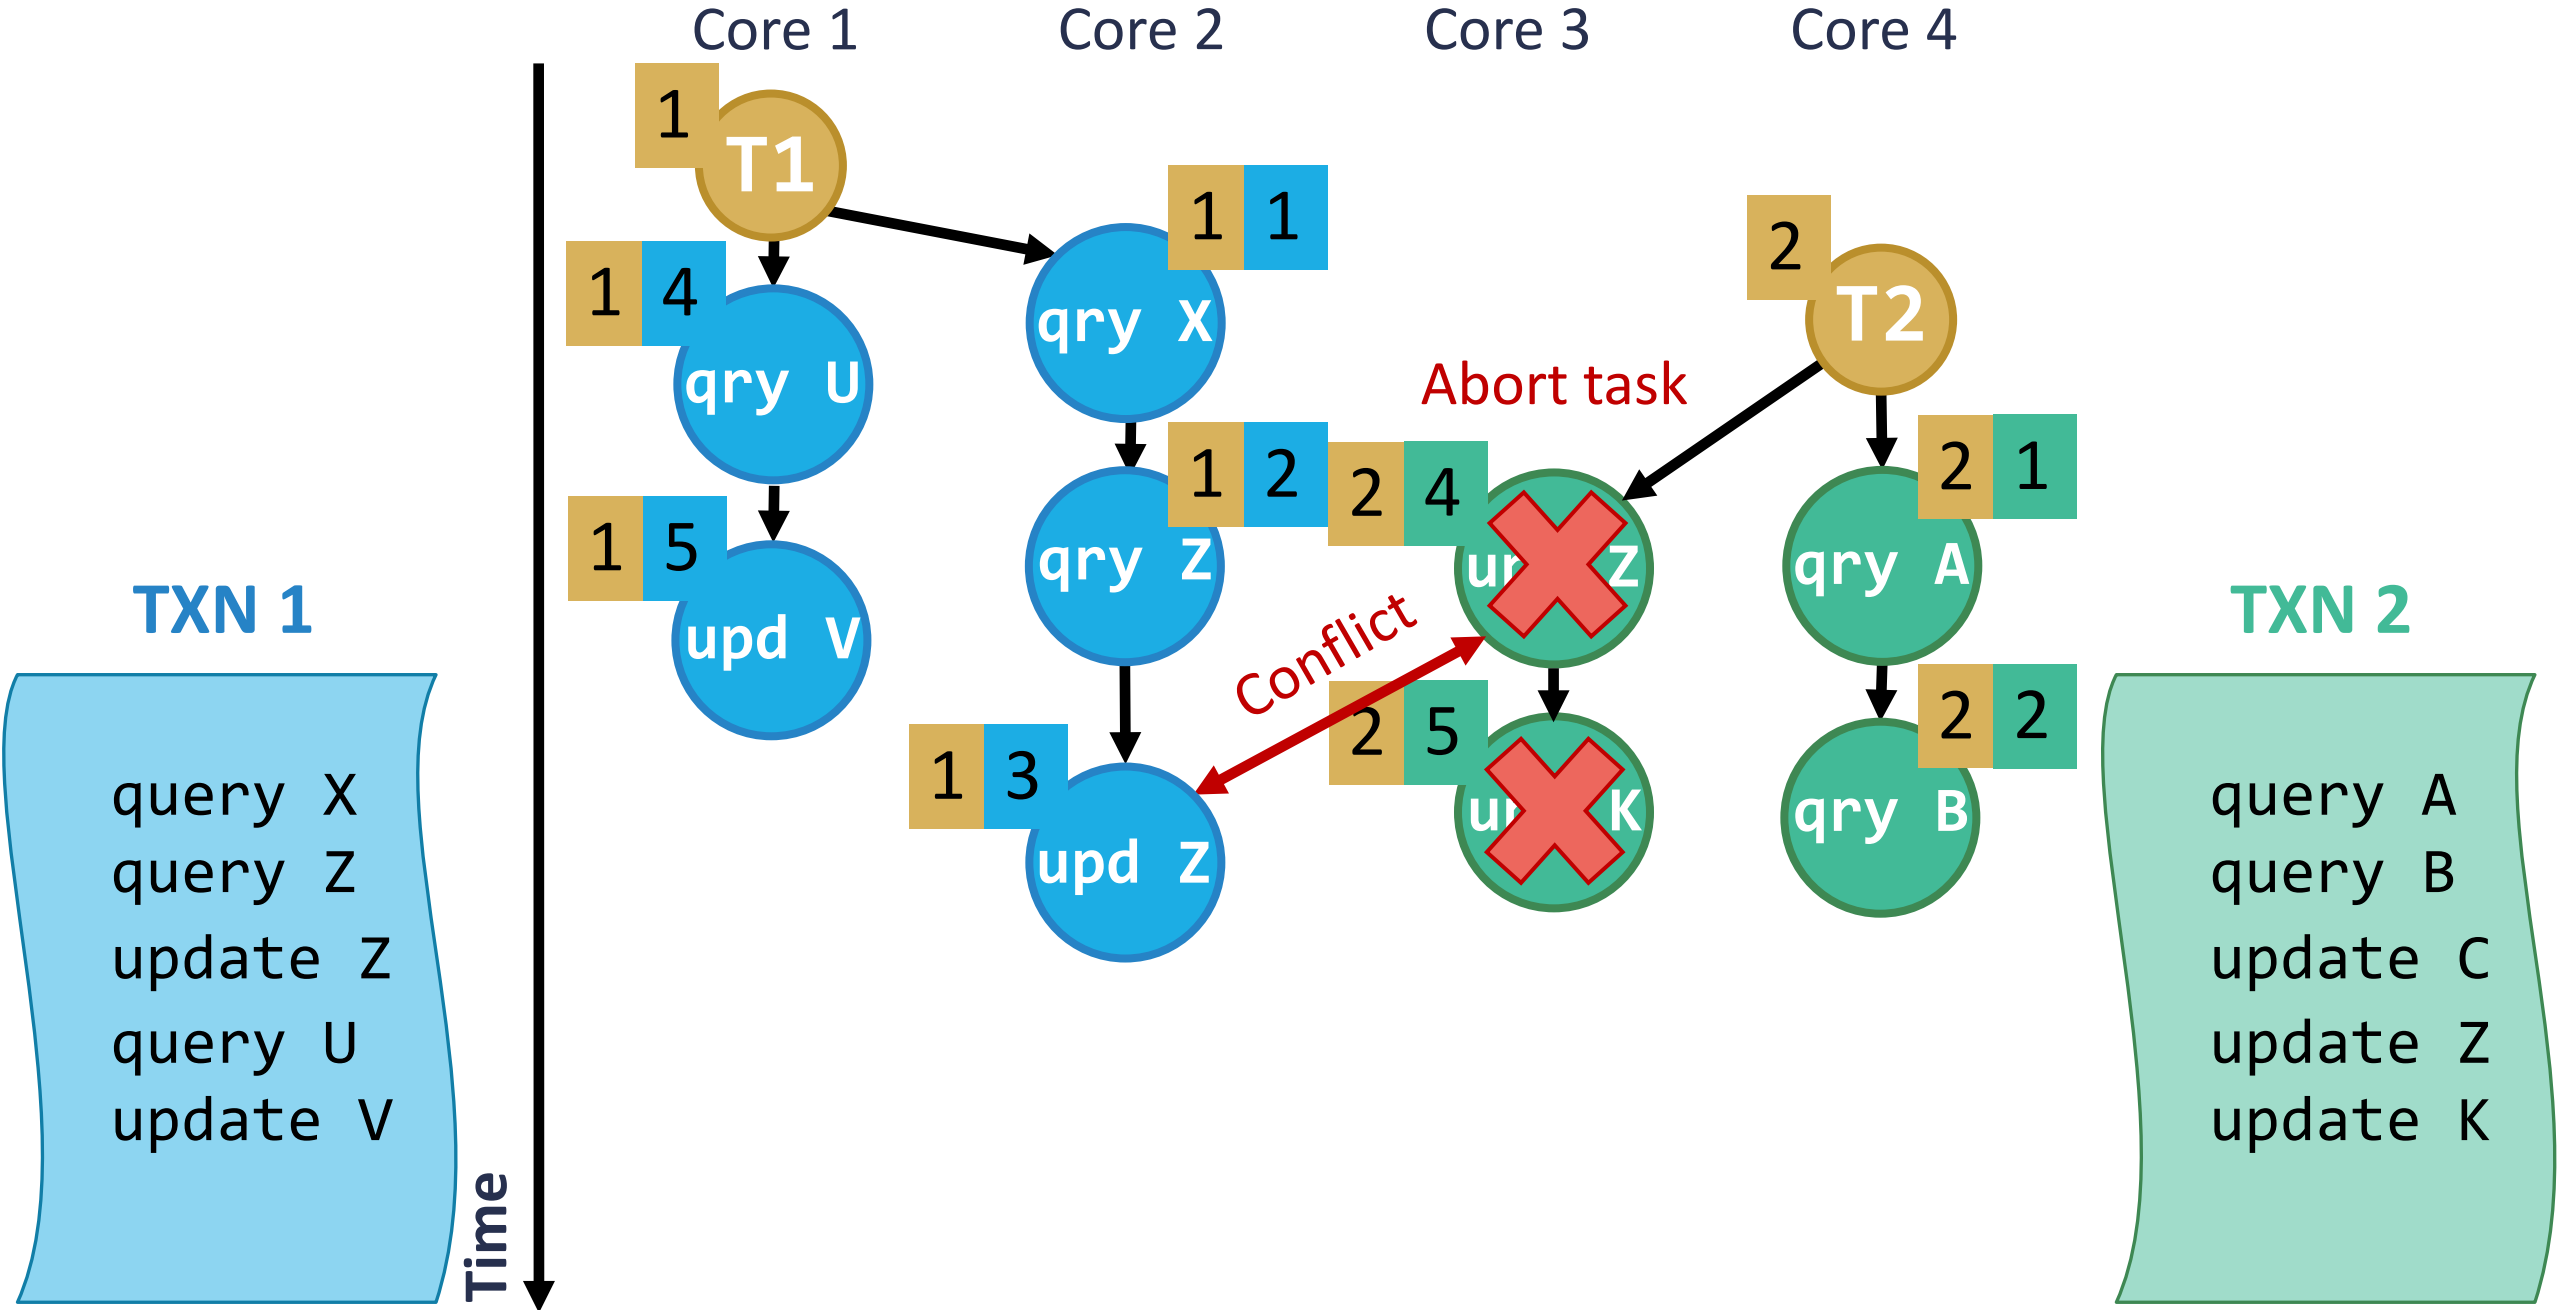


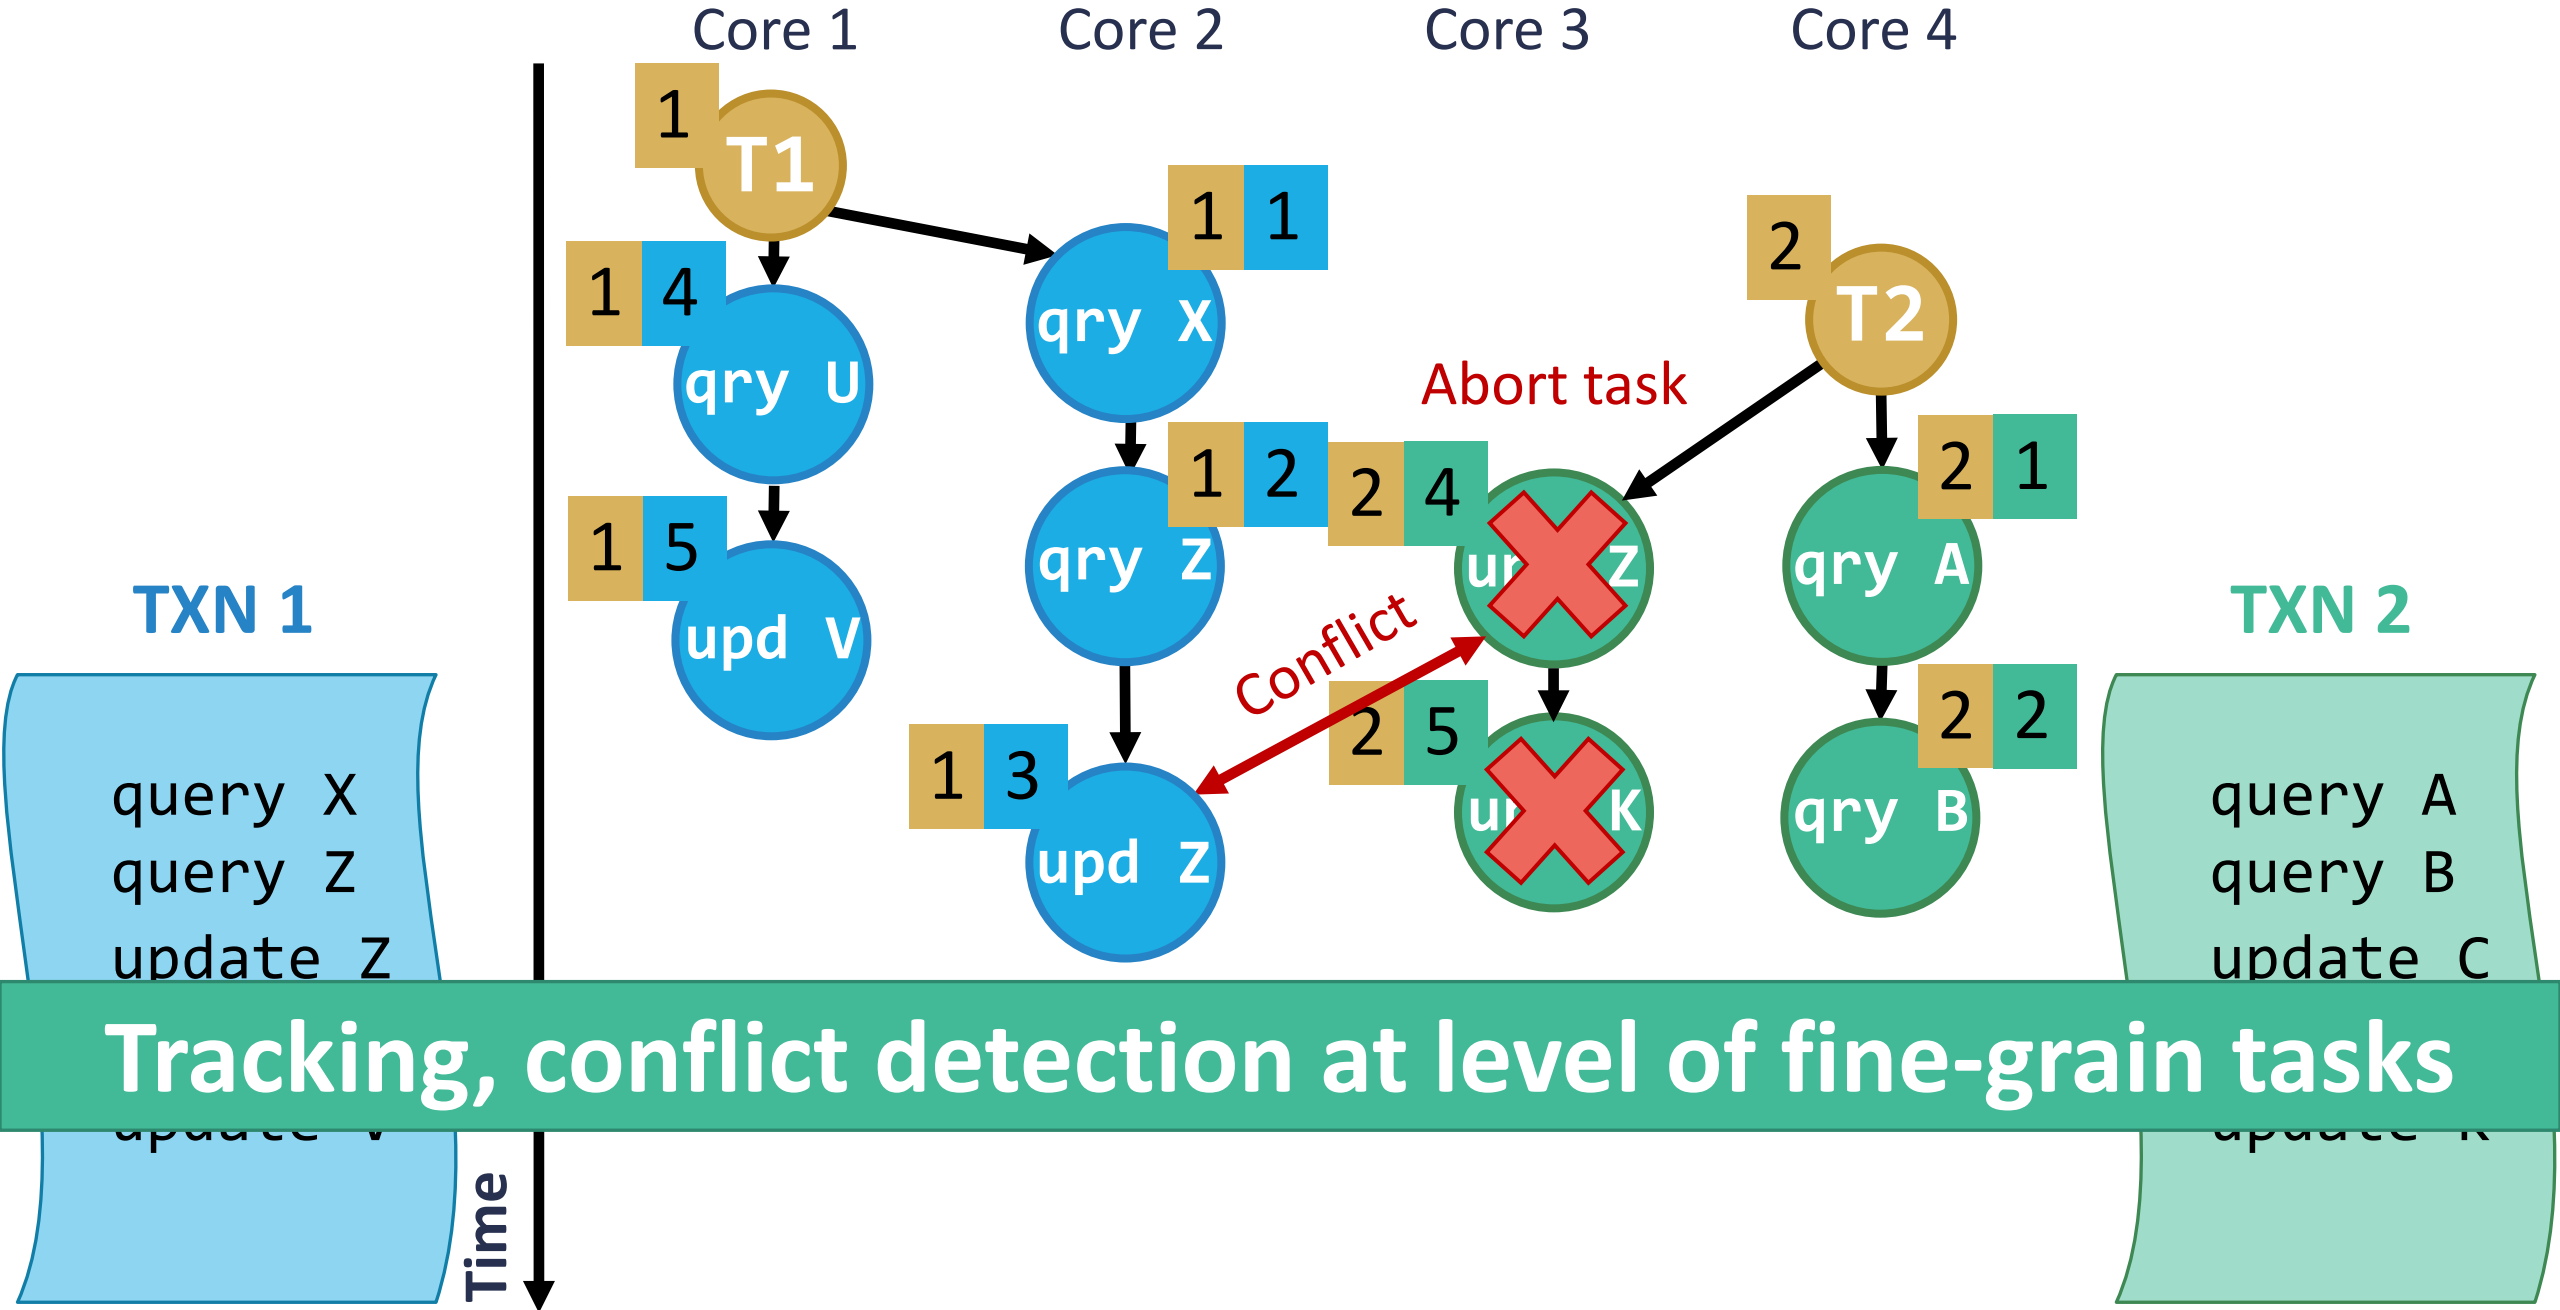


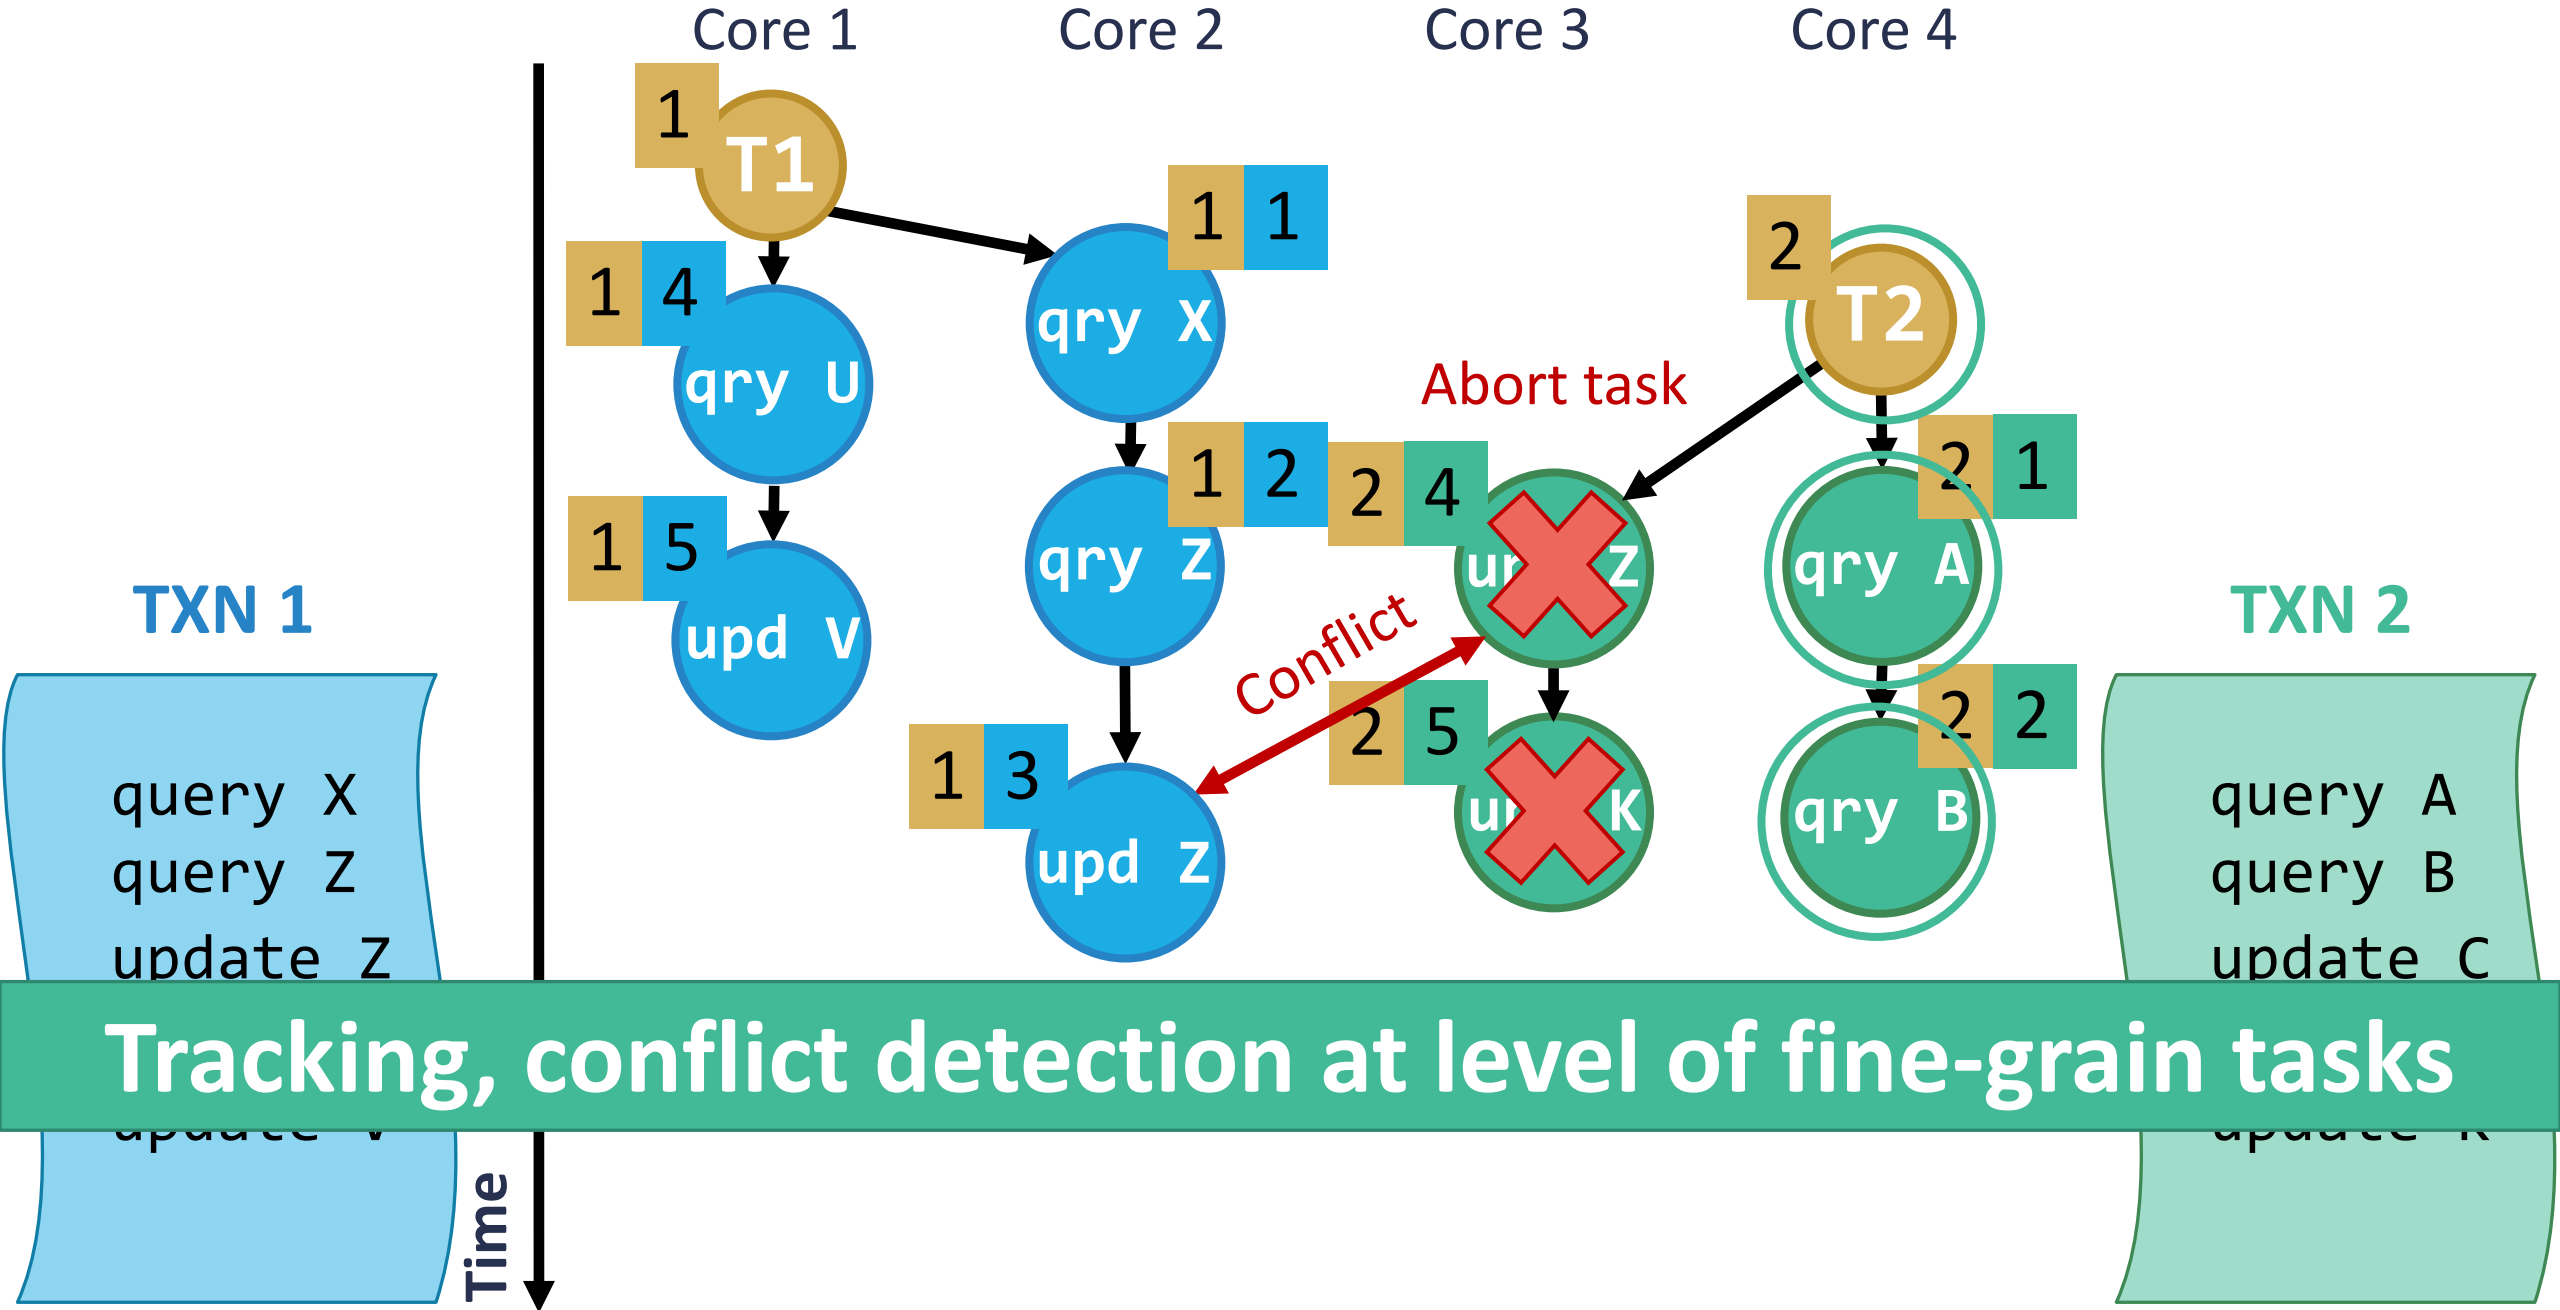


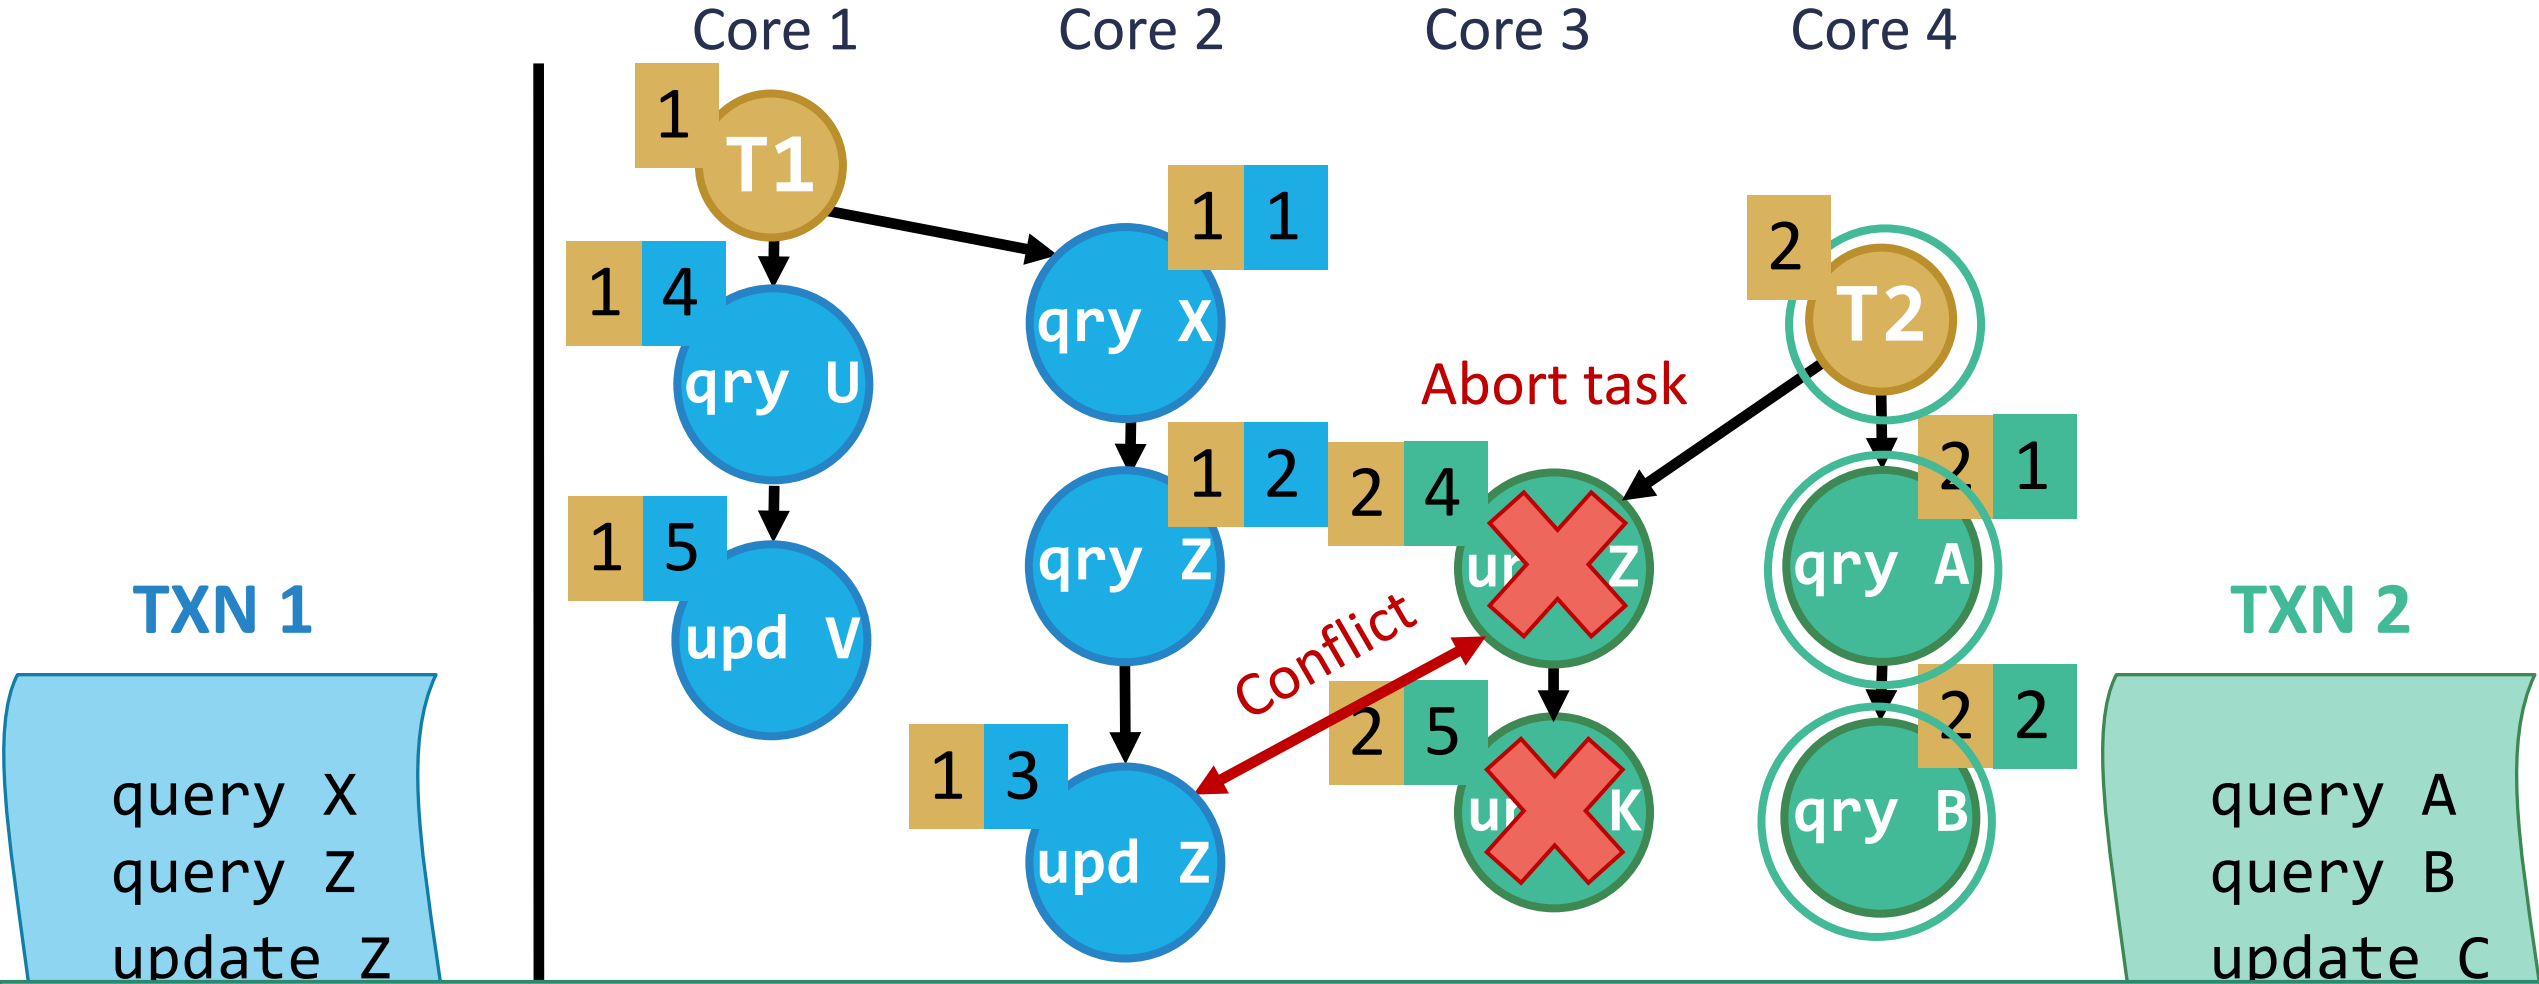








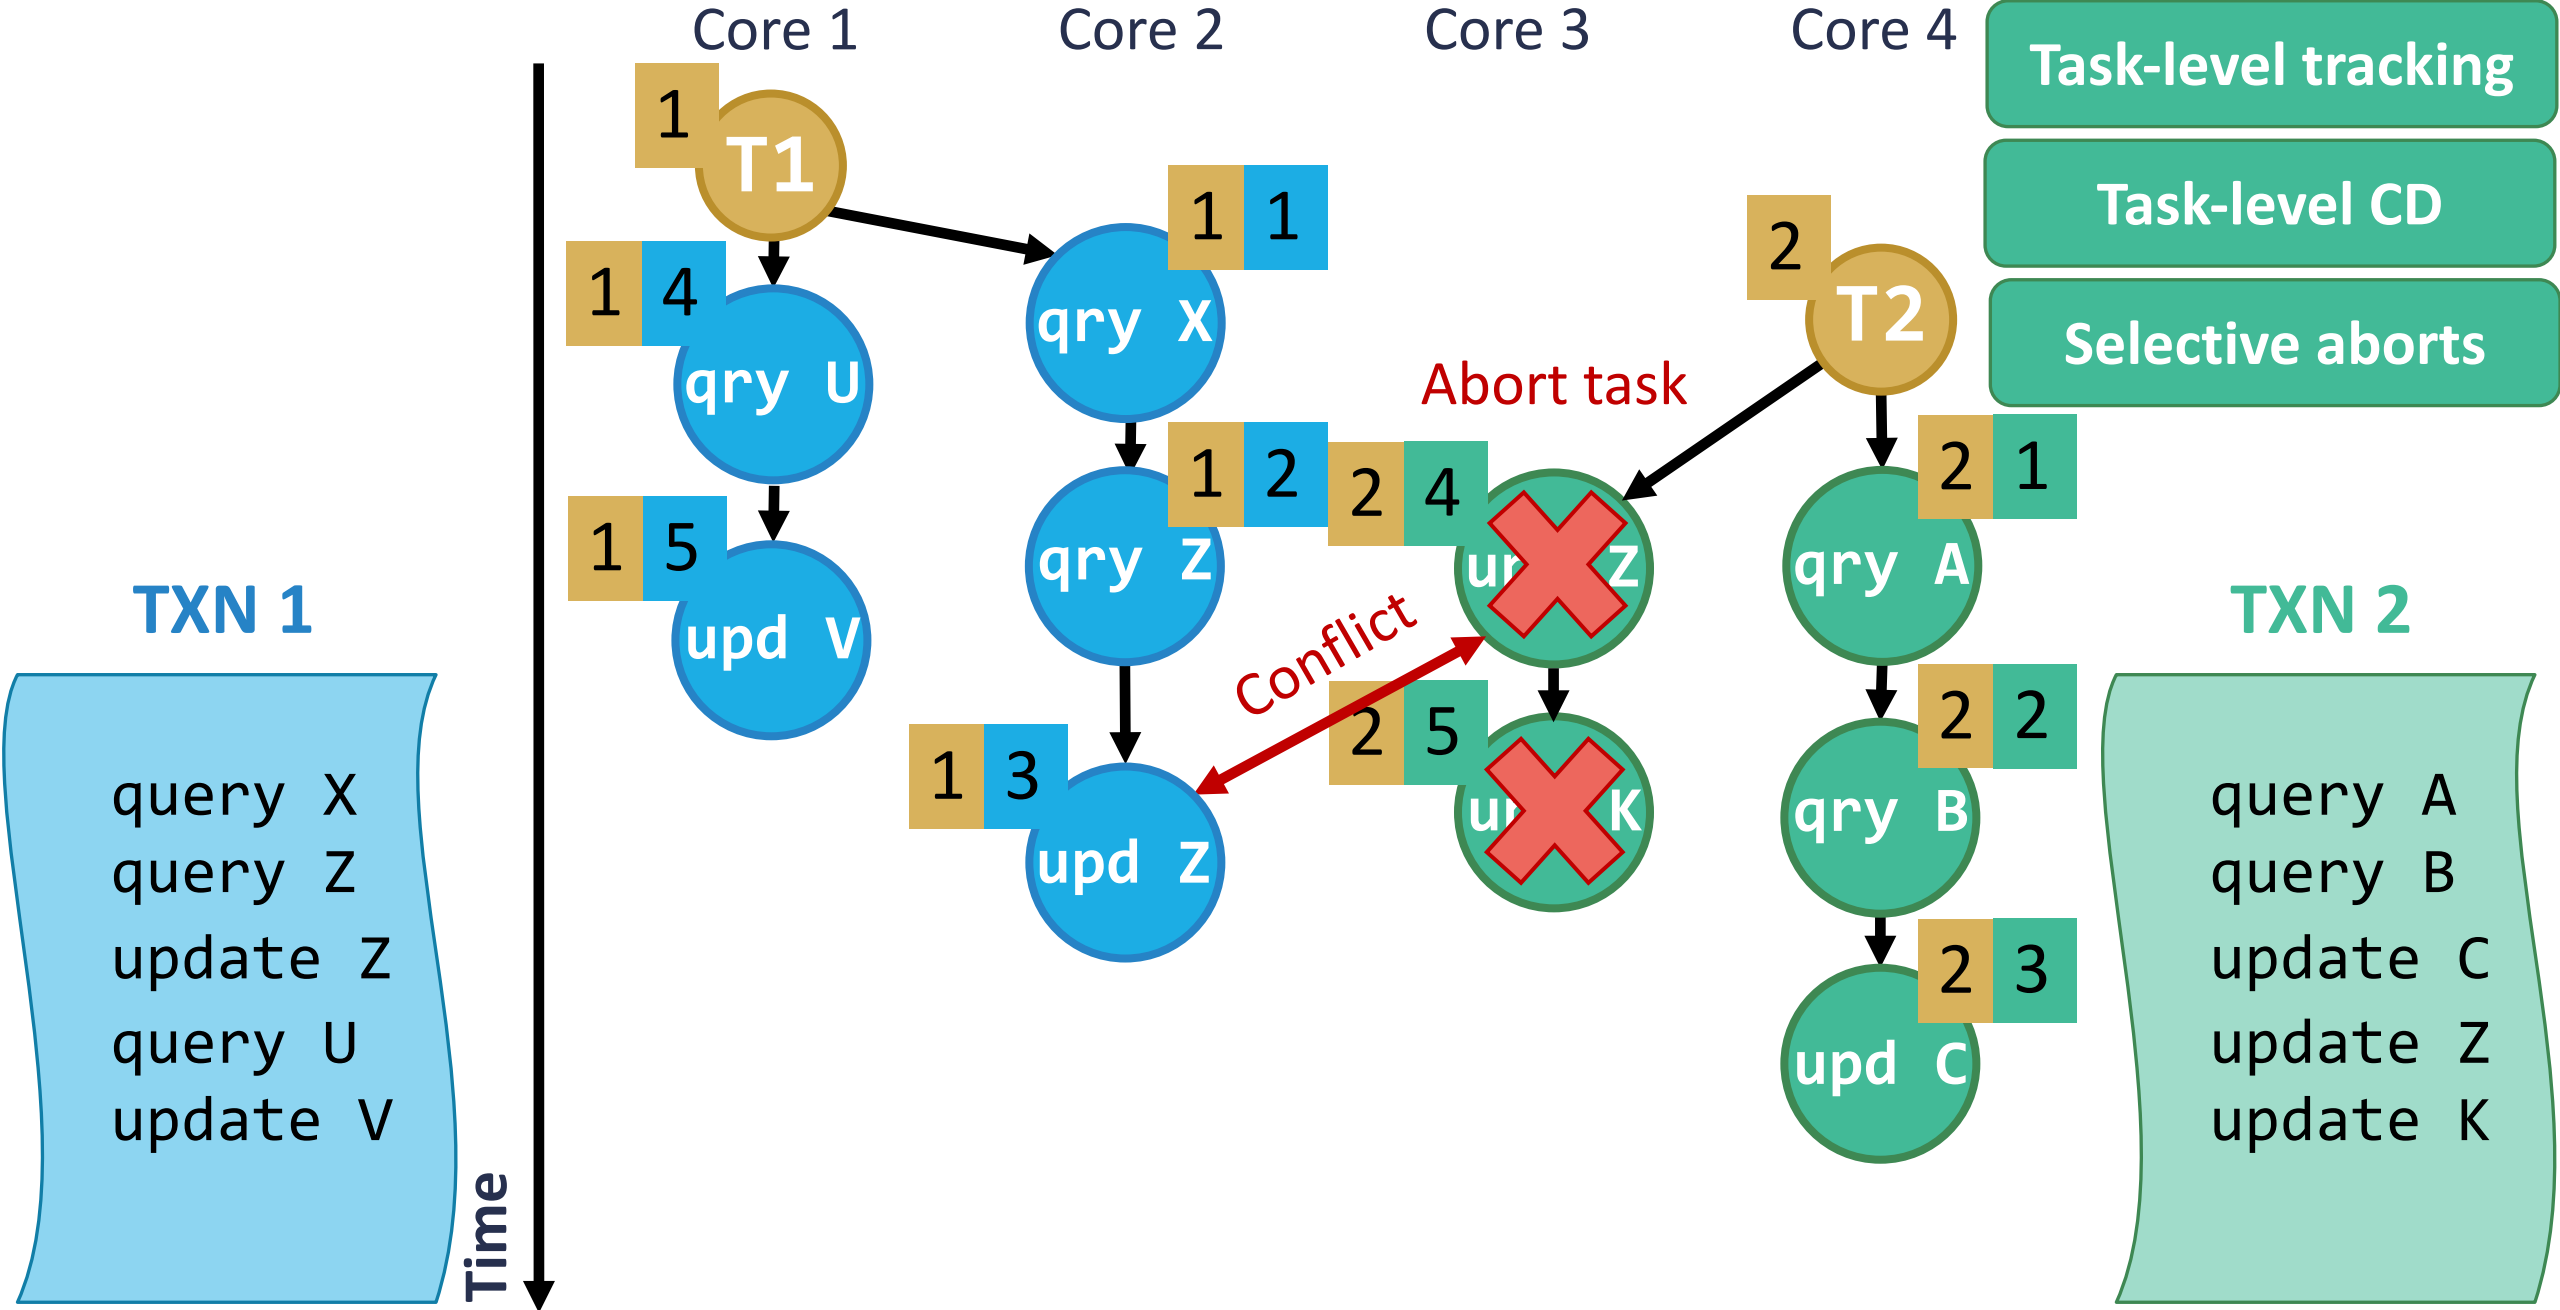


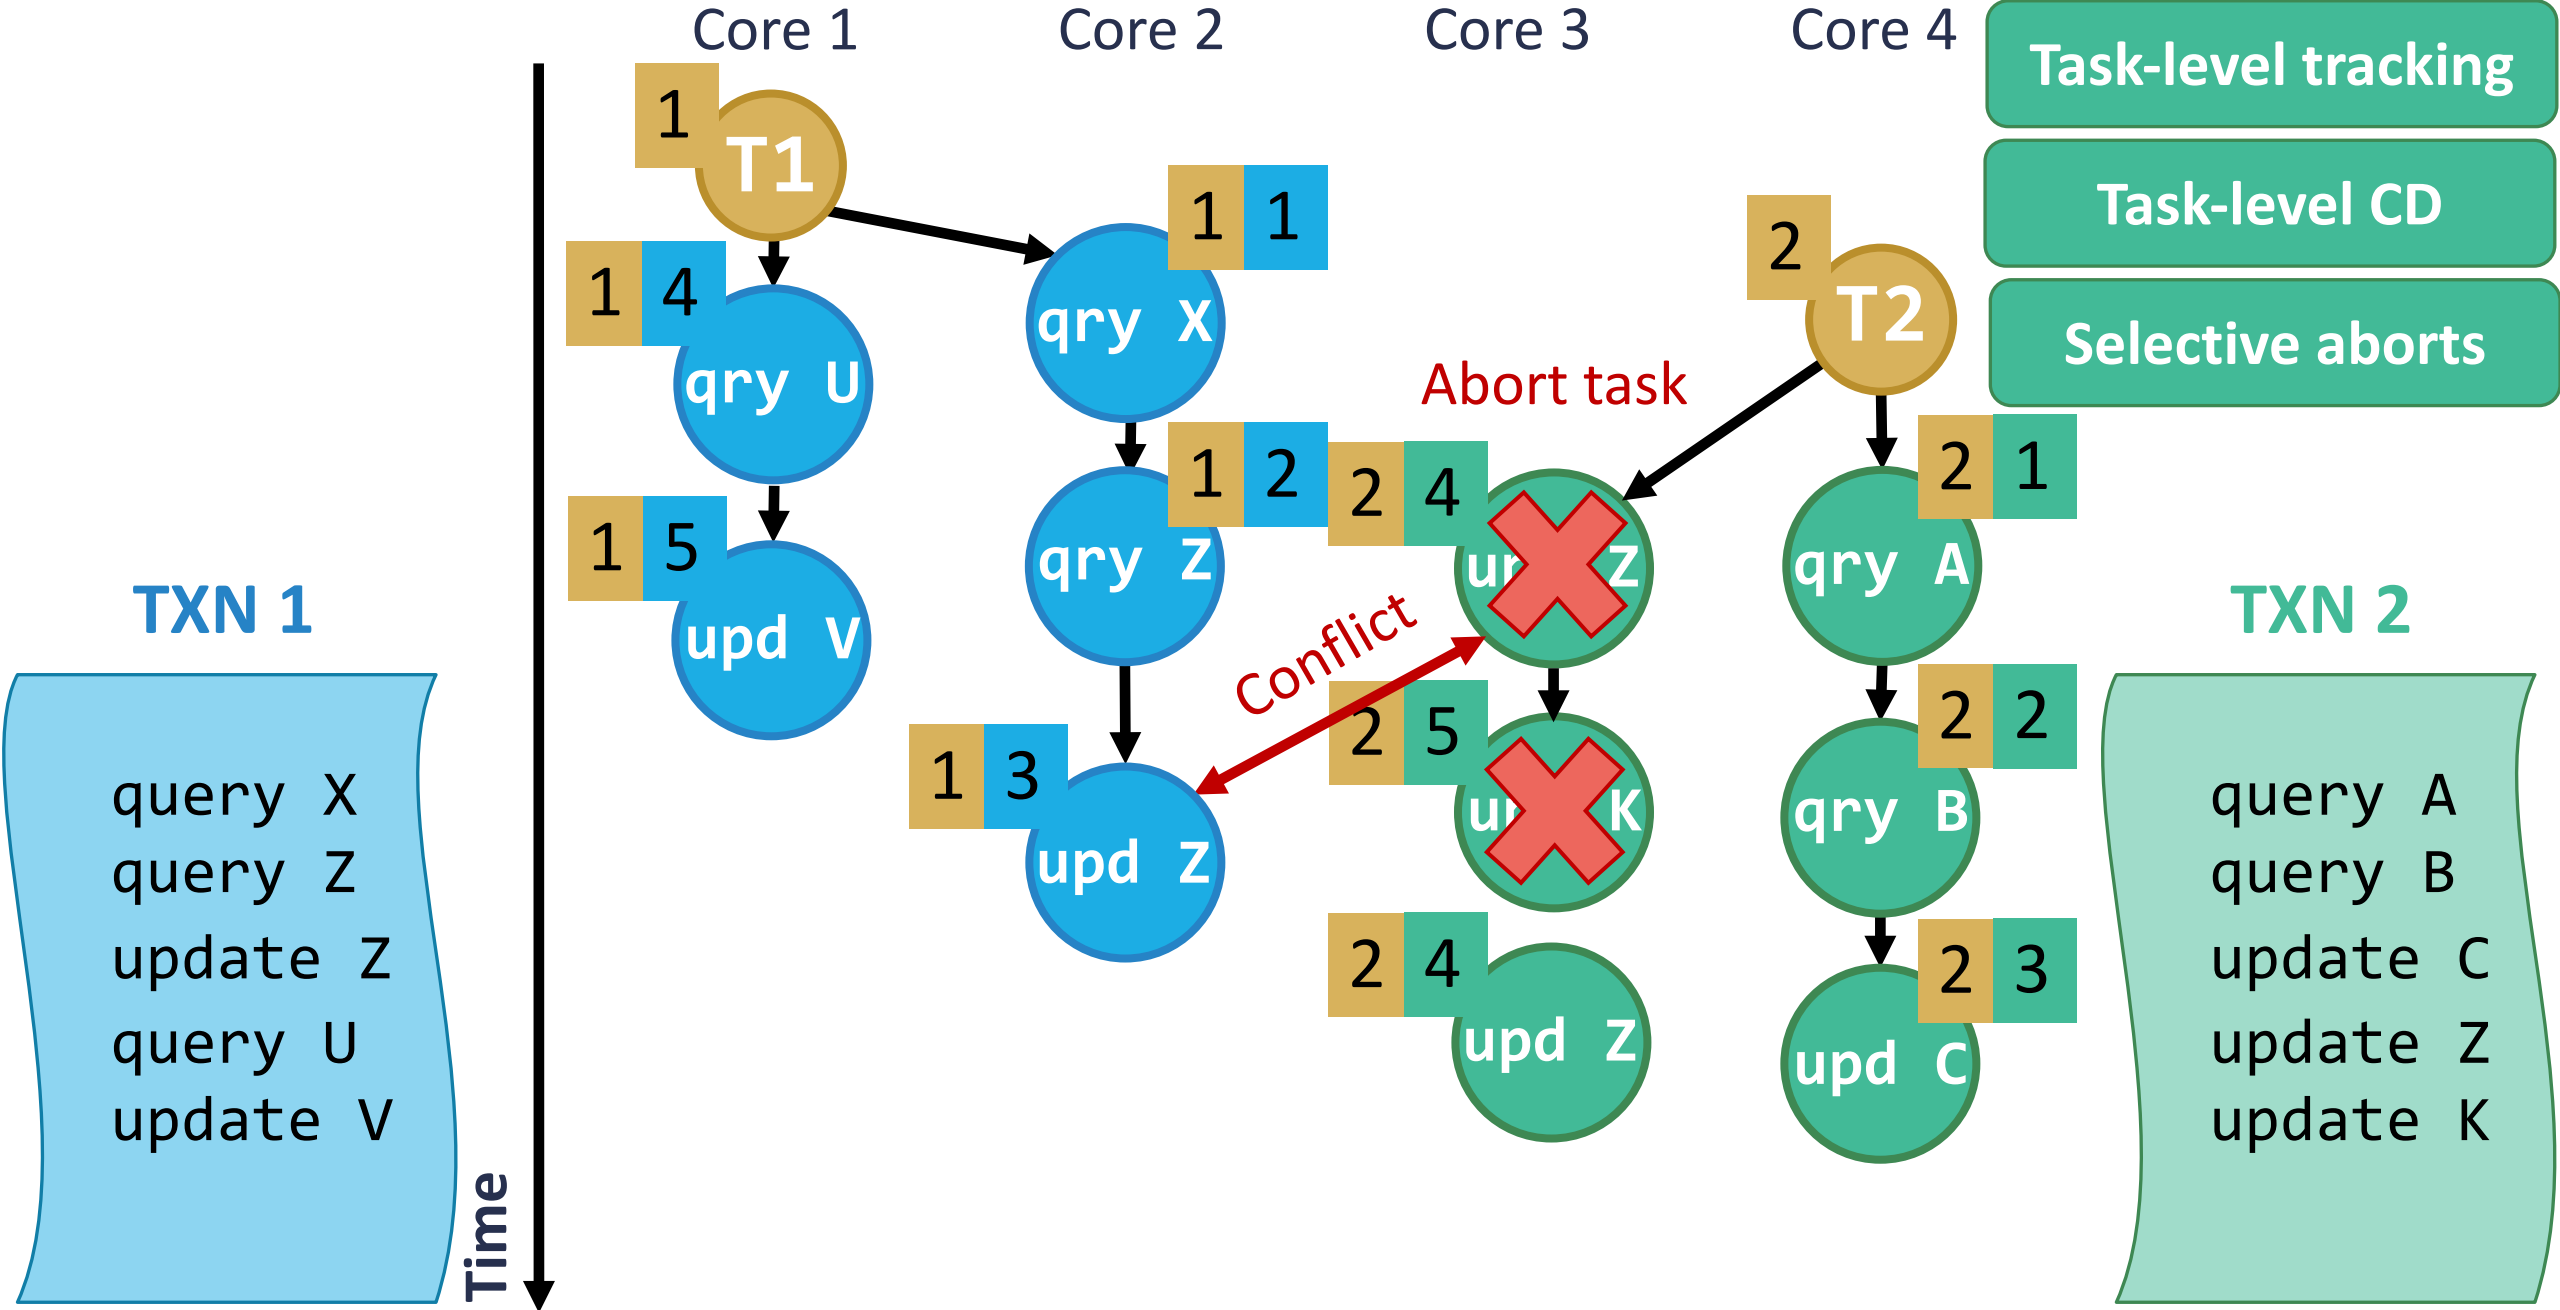


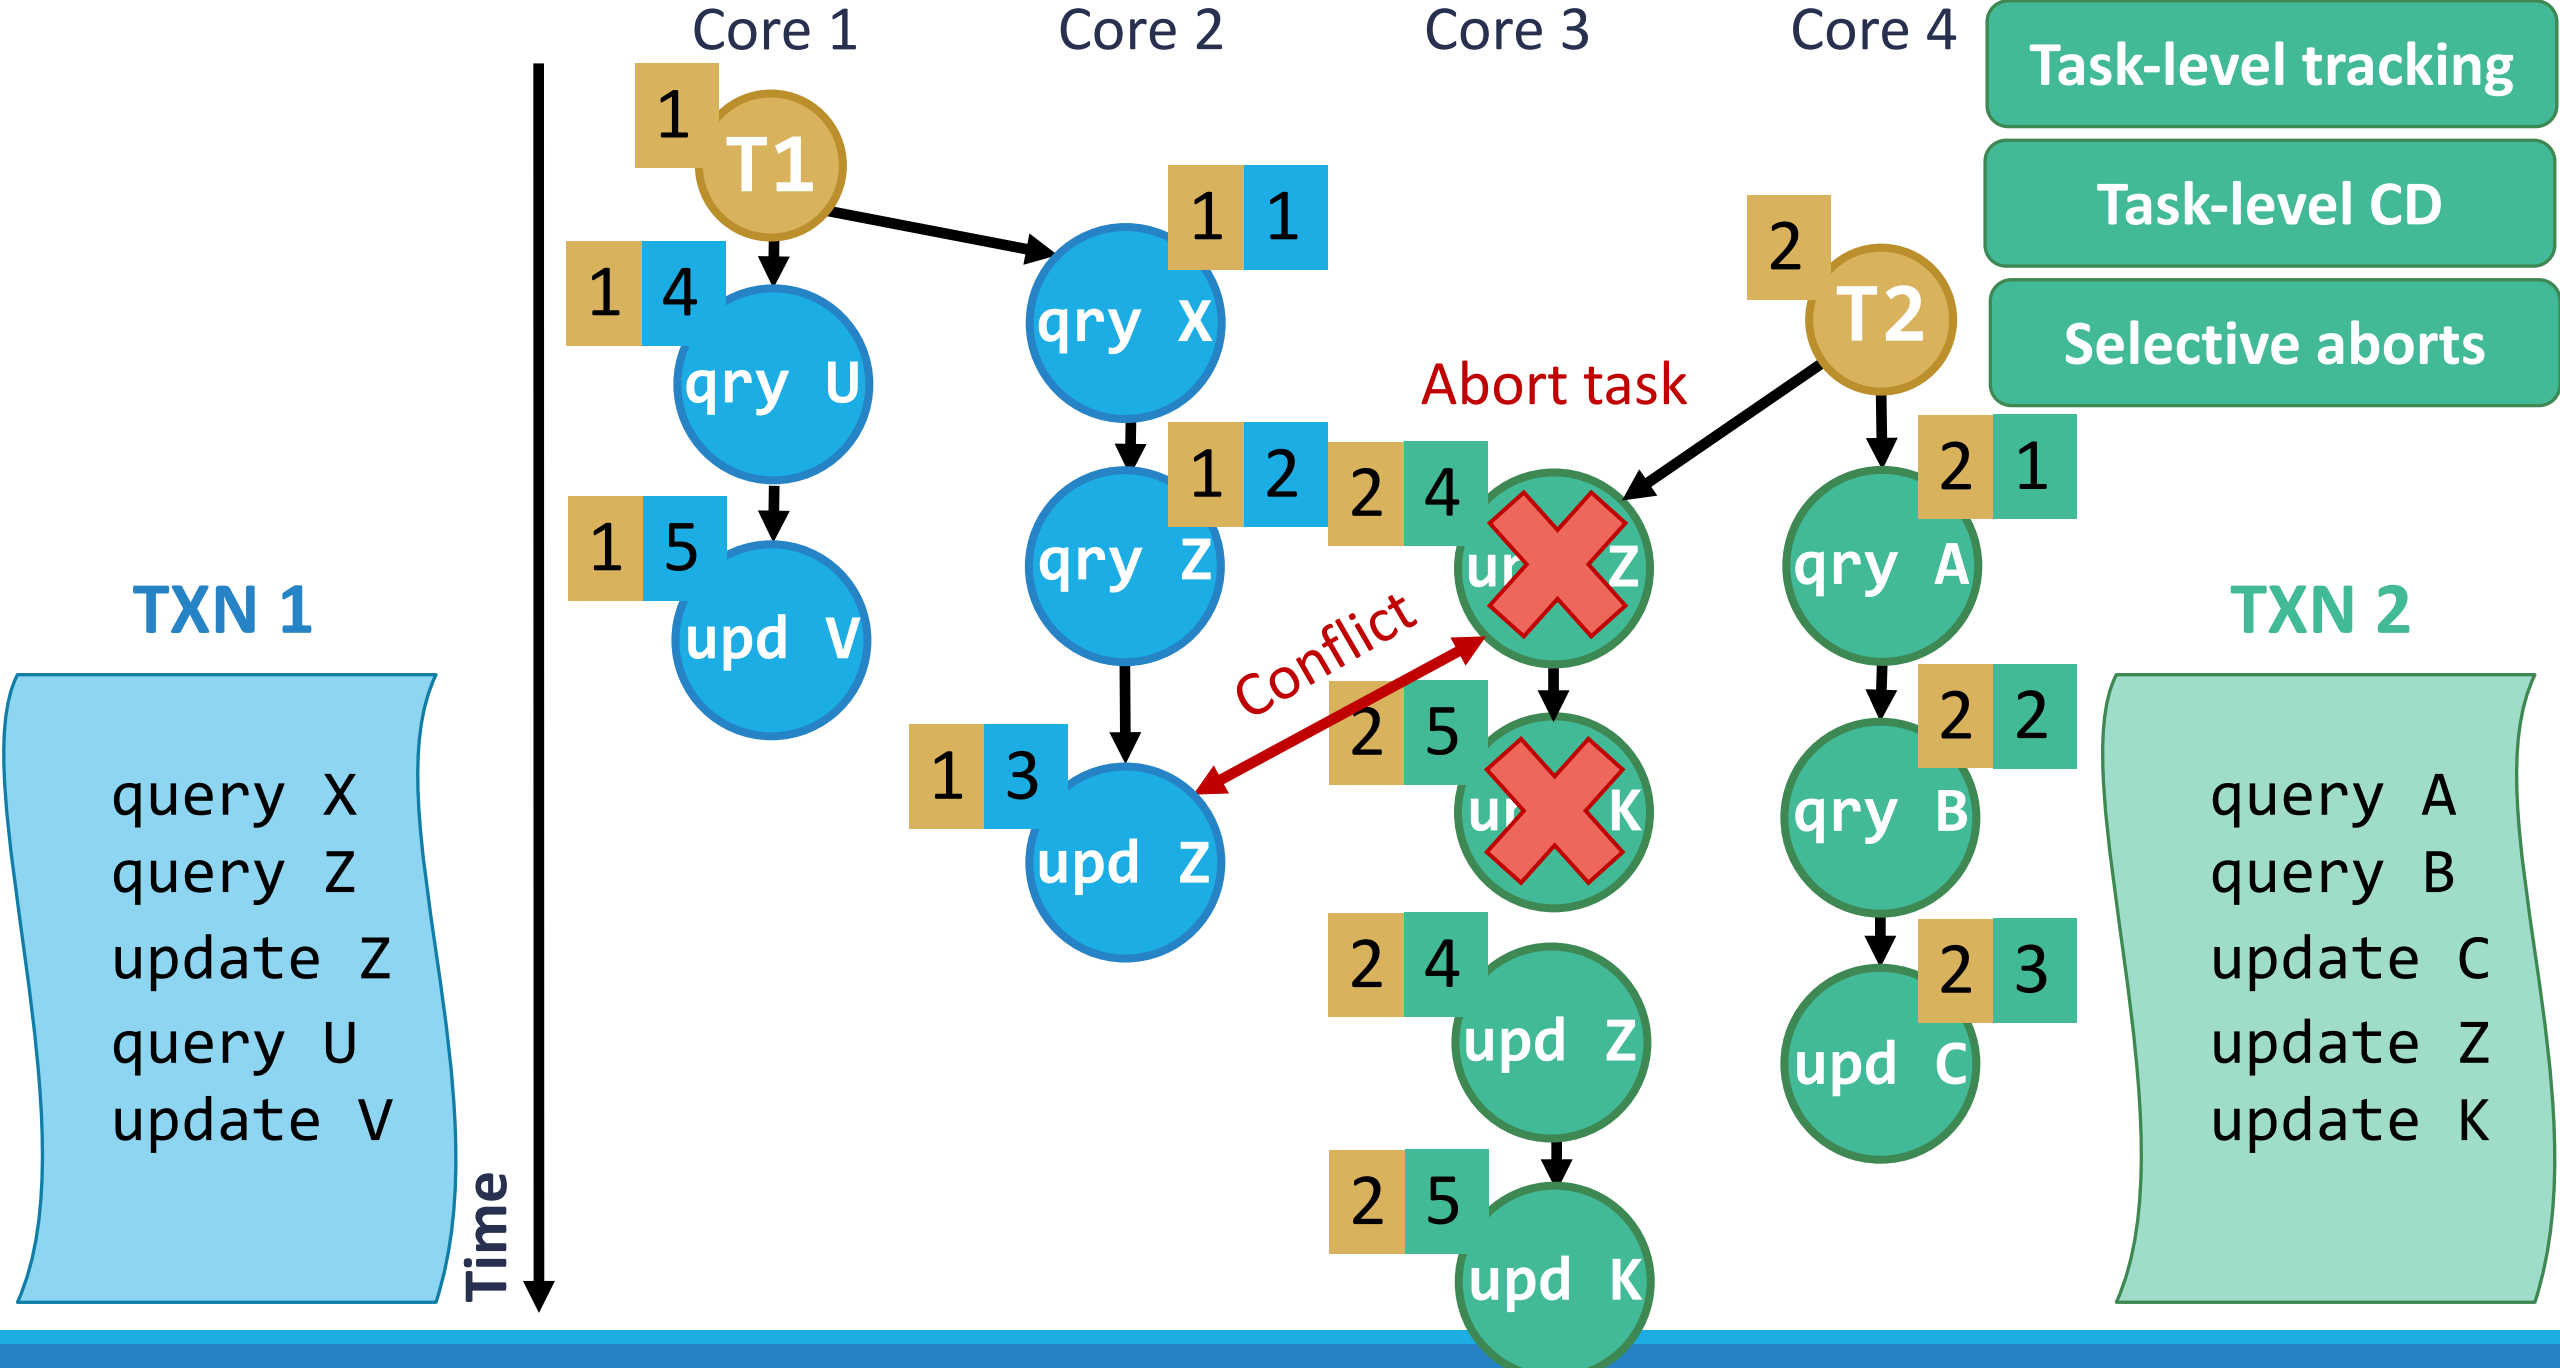
Tracking, conflict detection at level of fine-grain tasks

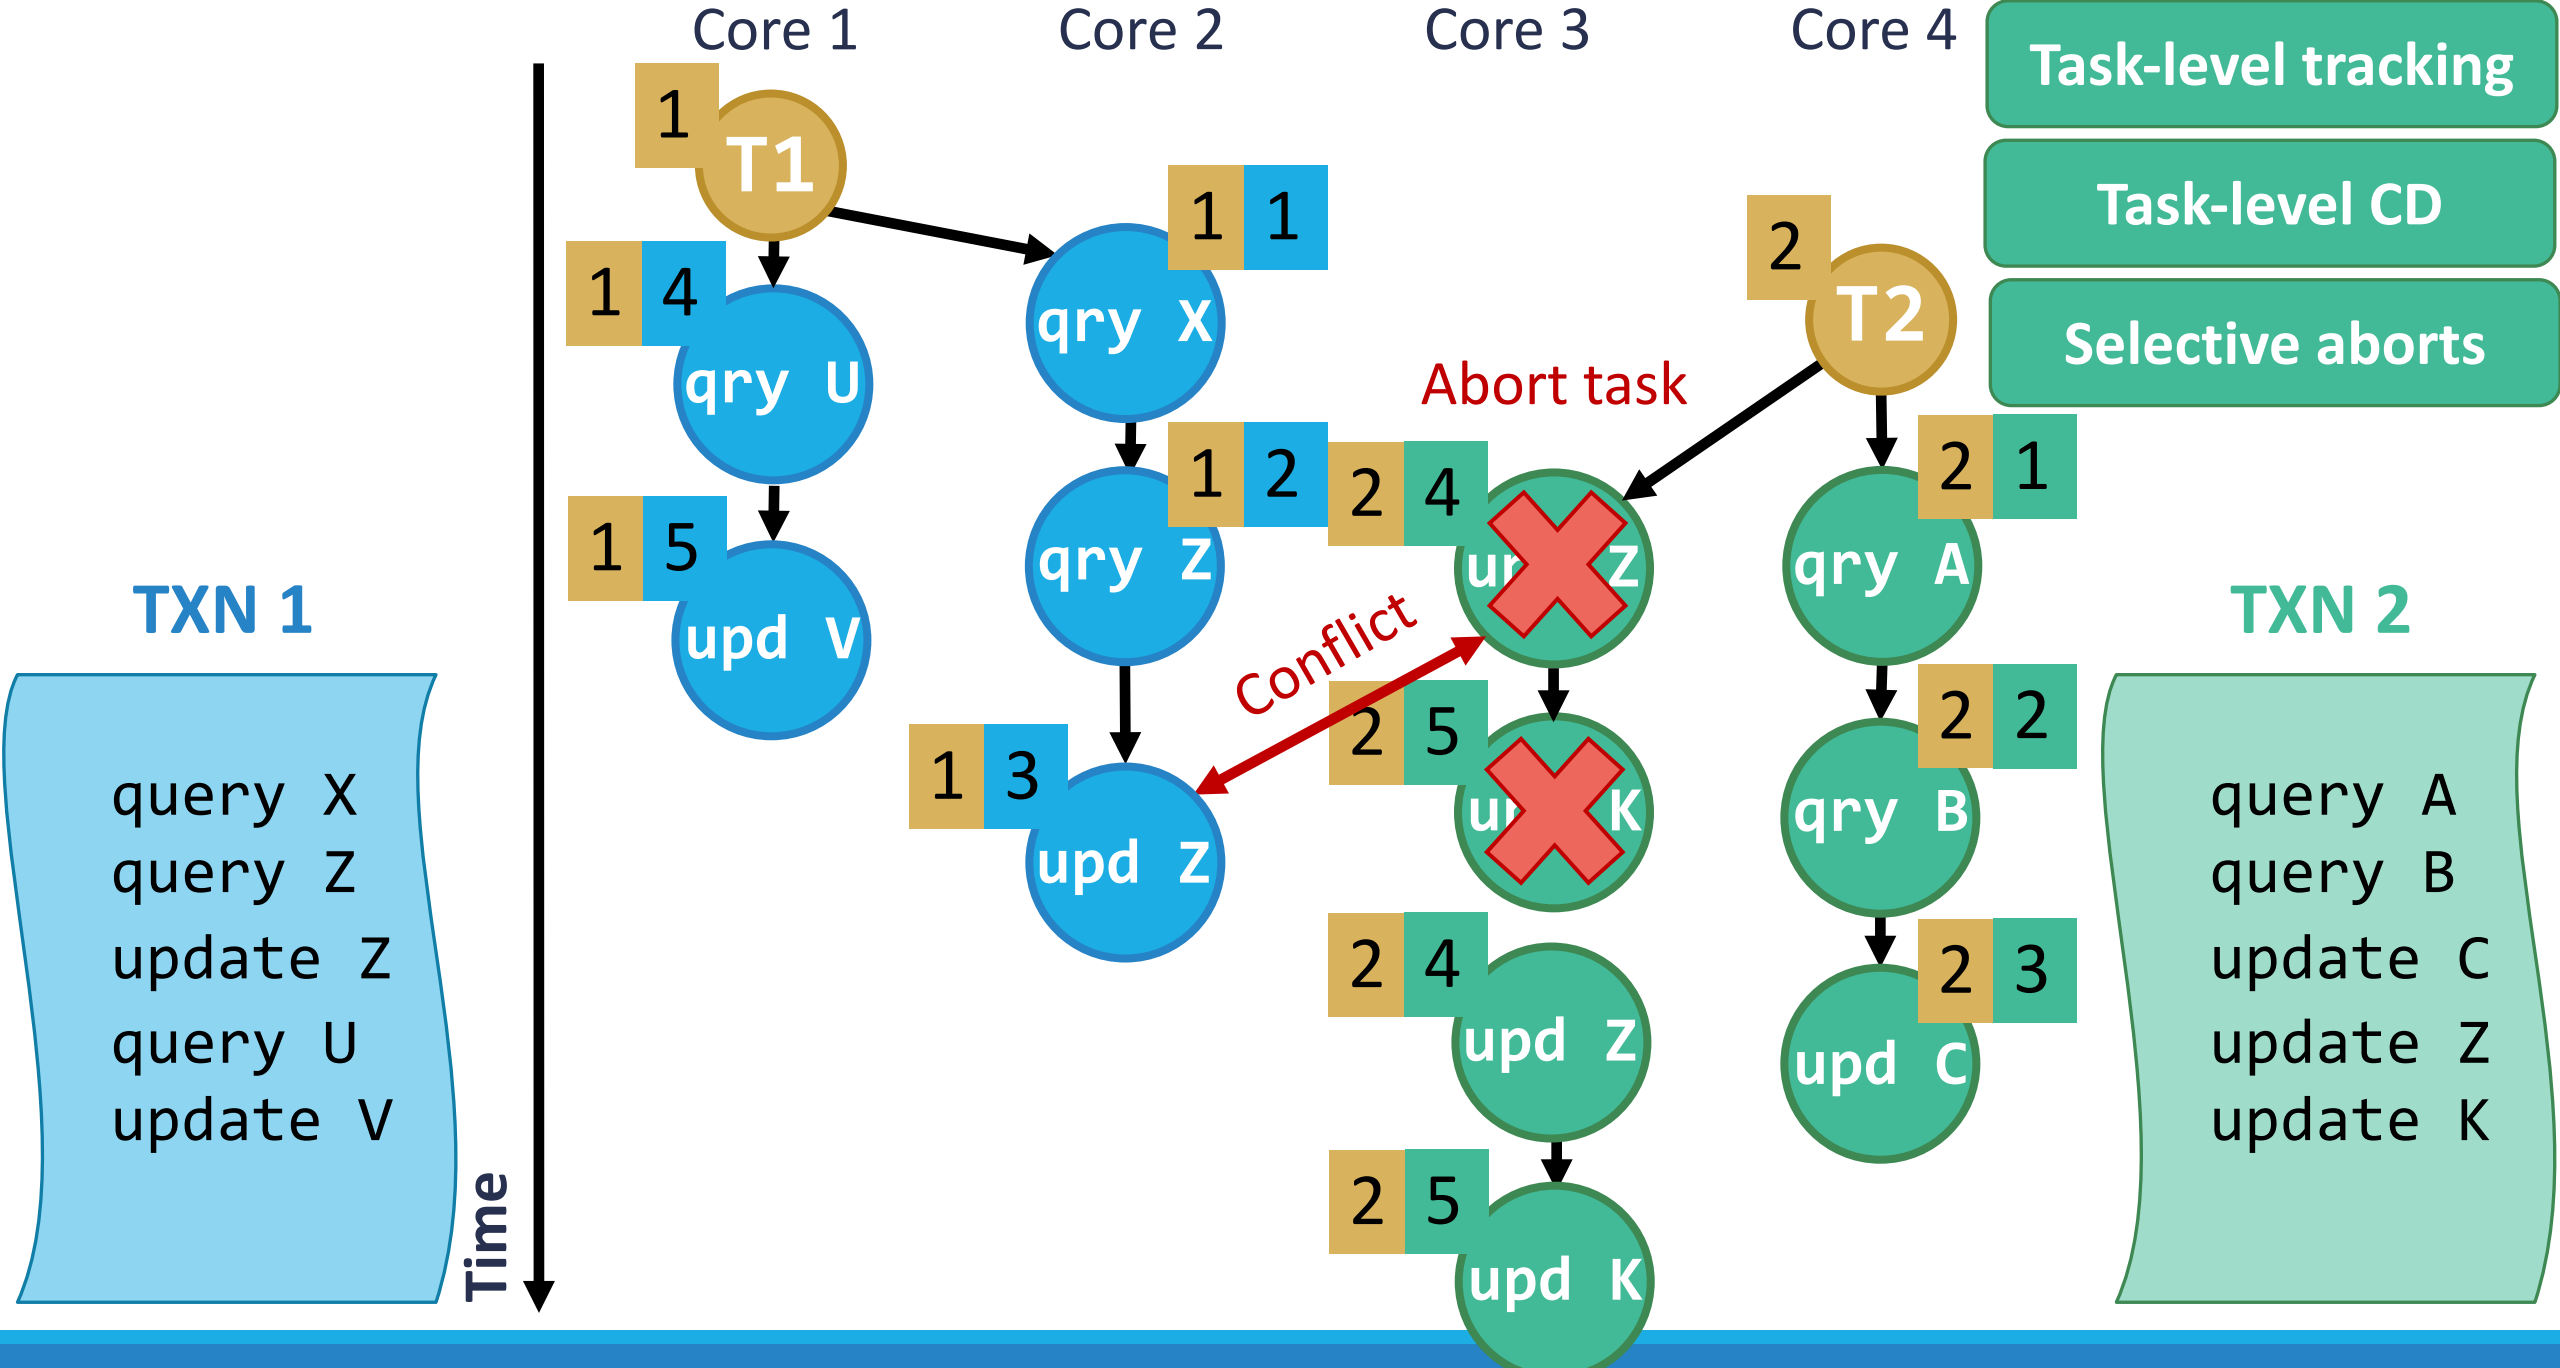
Selective aborts waste less work

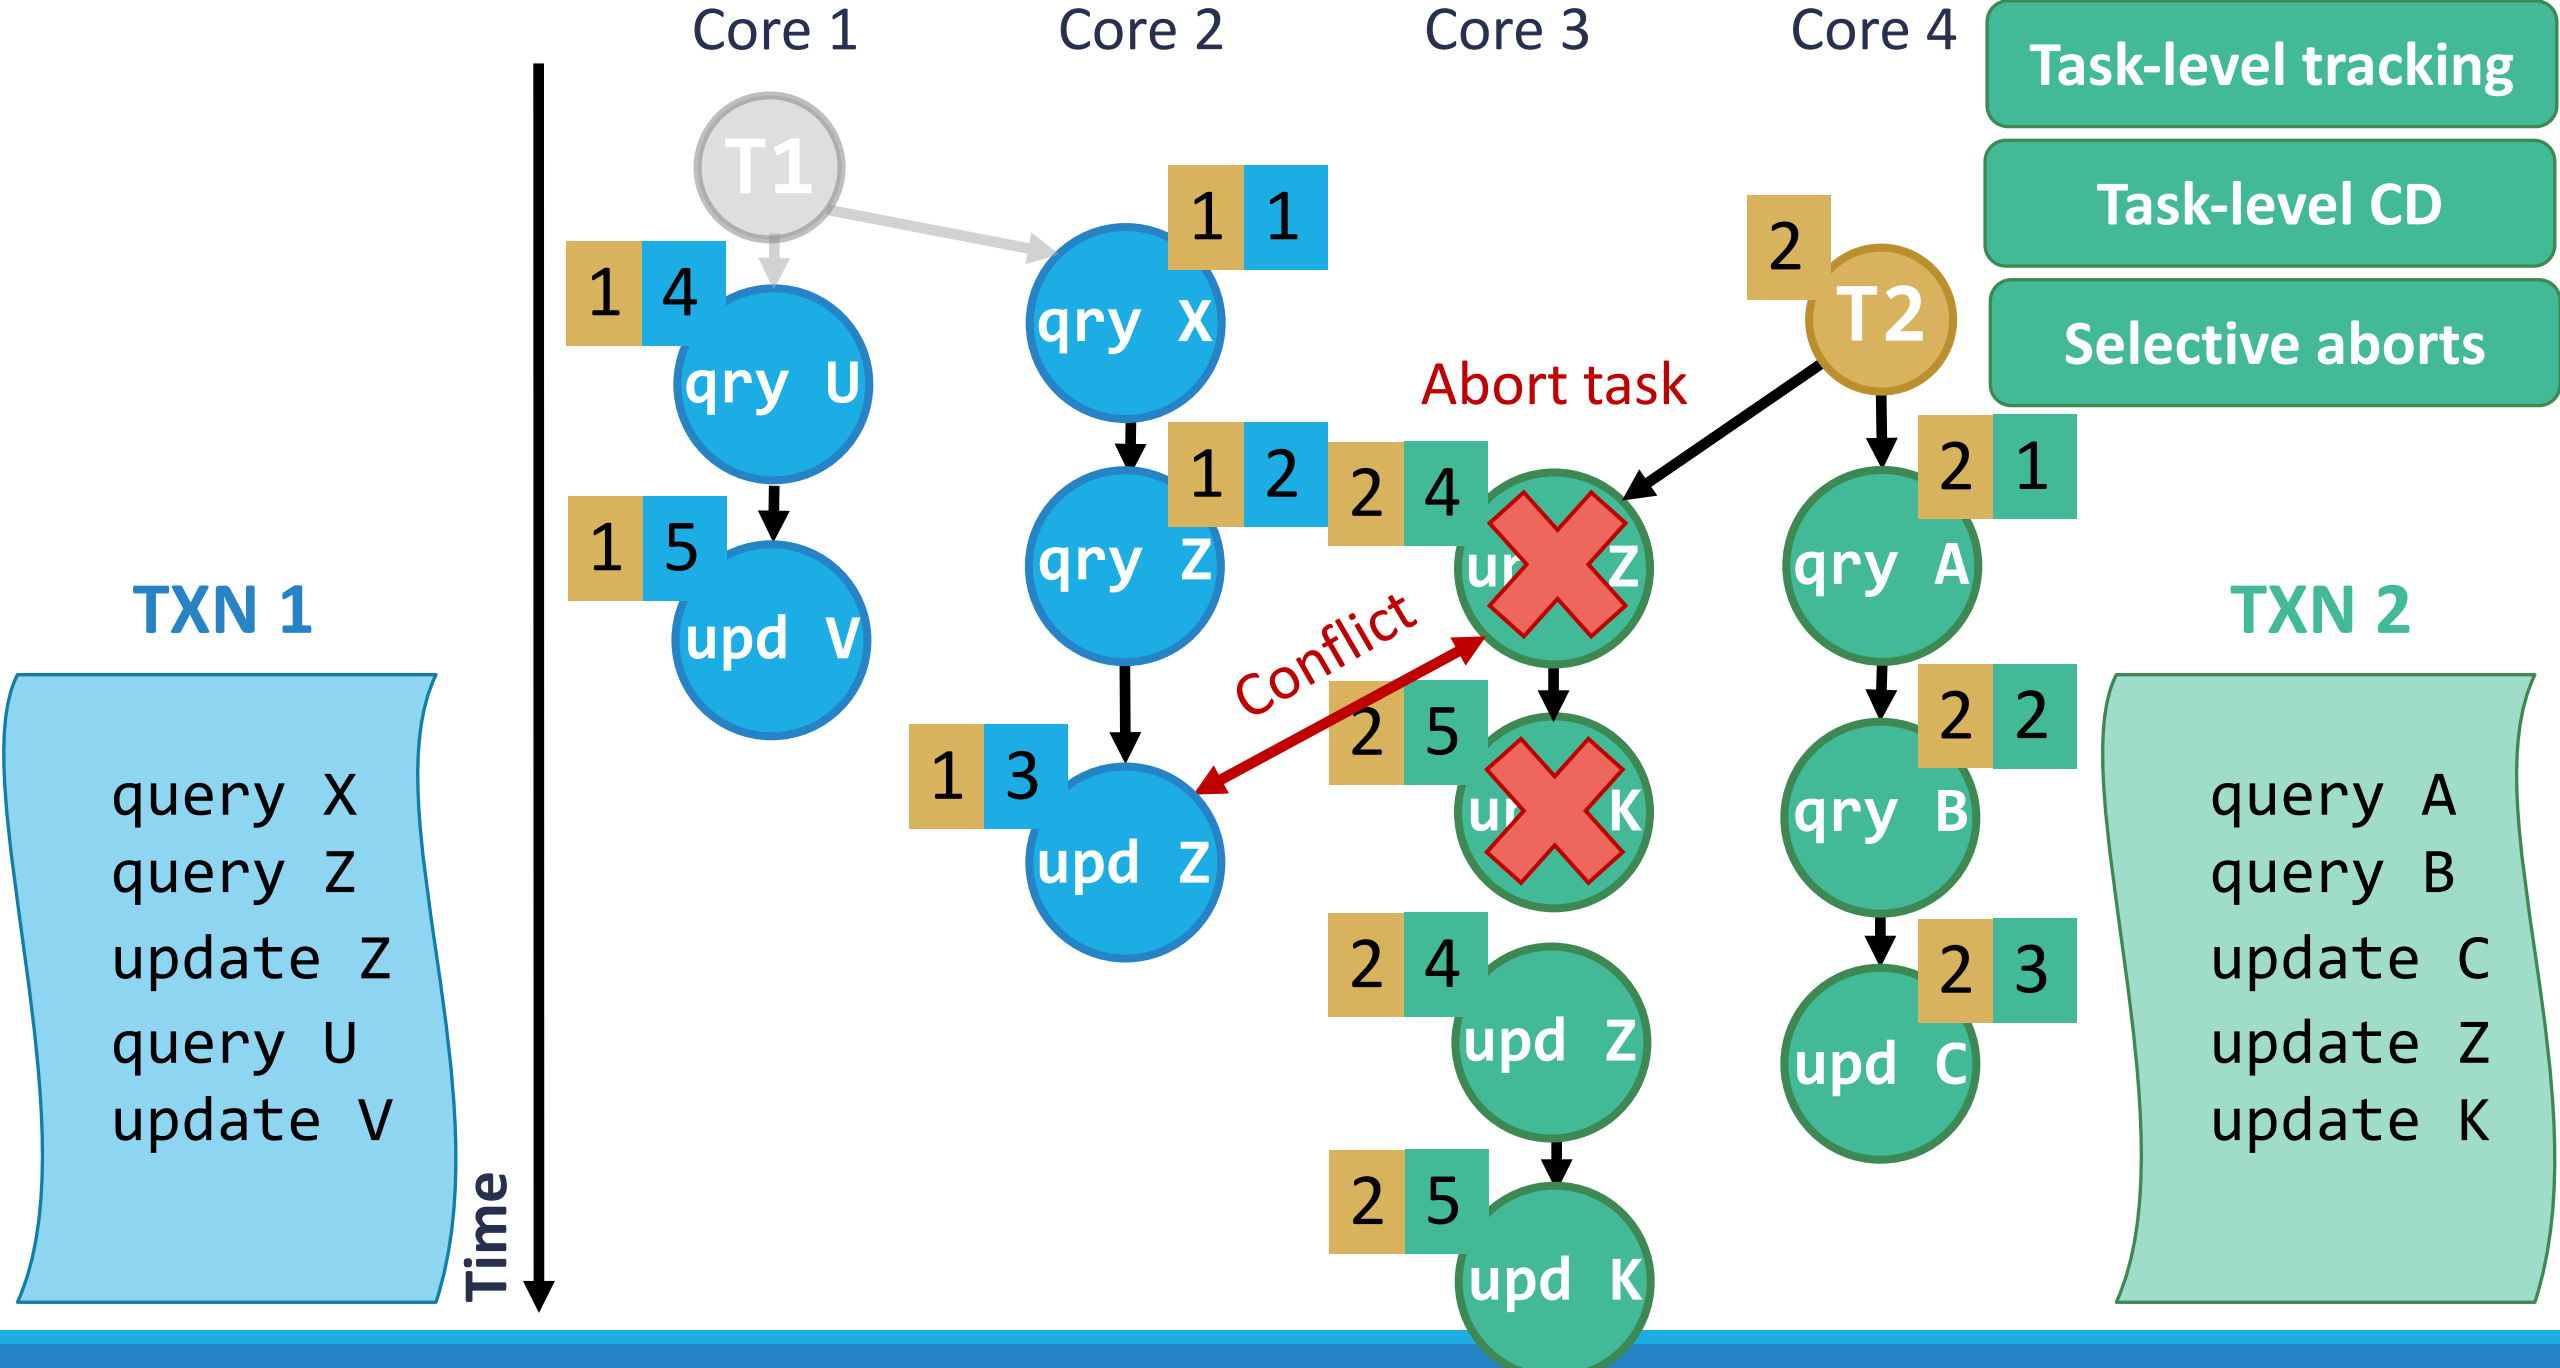


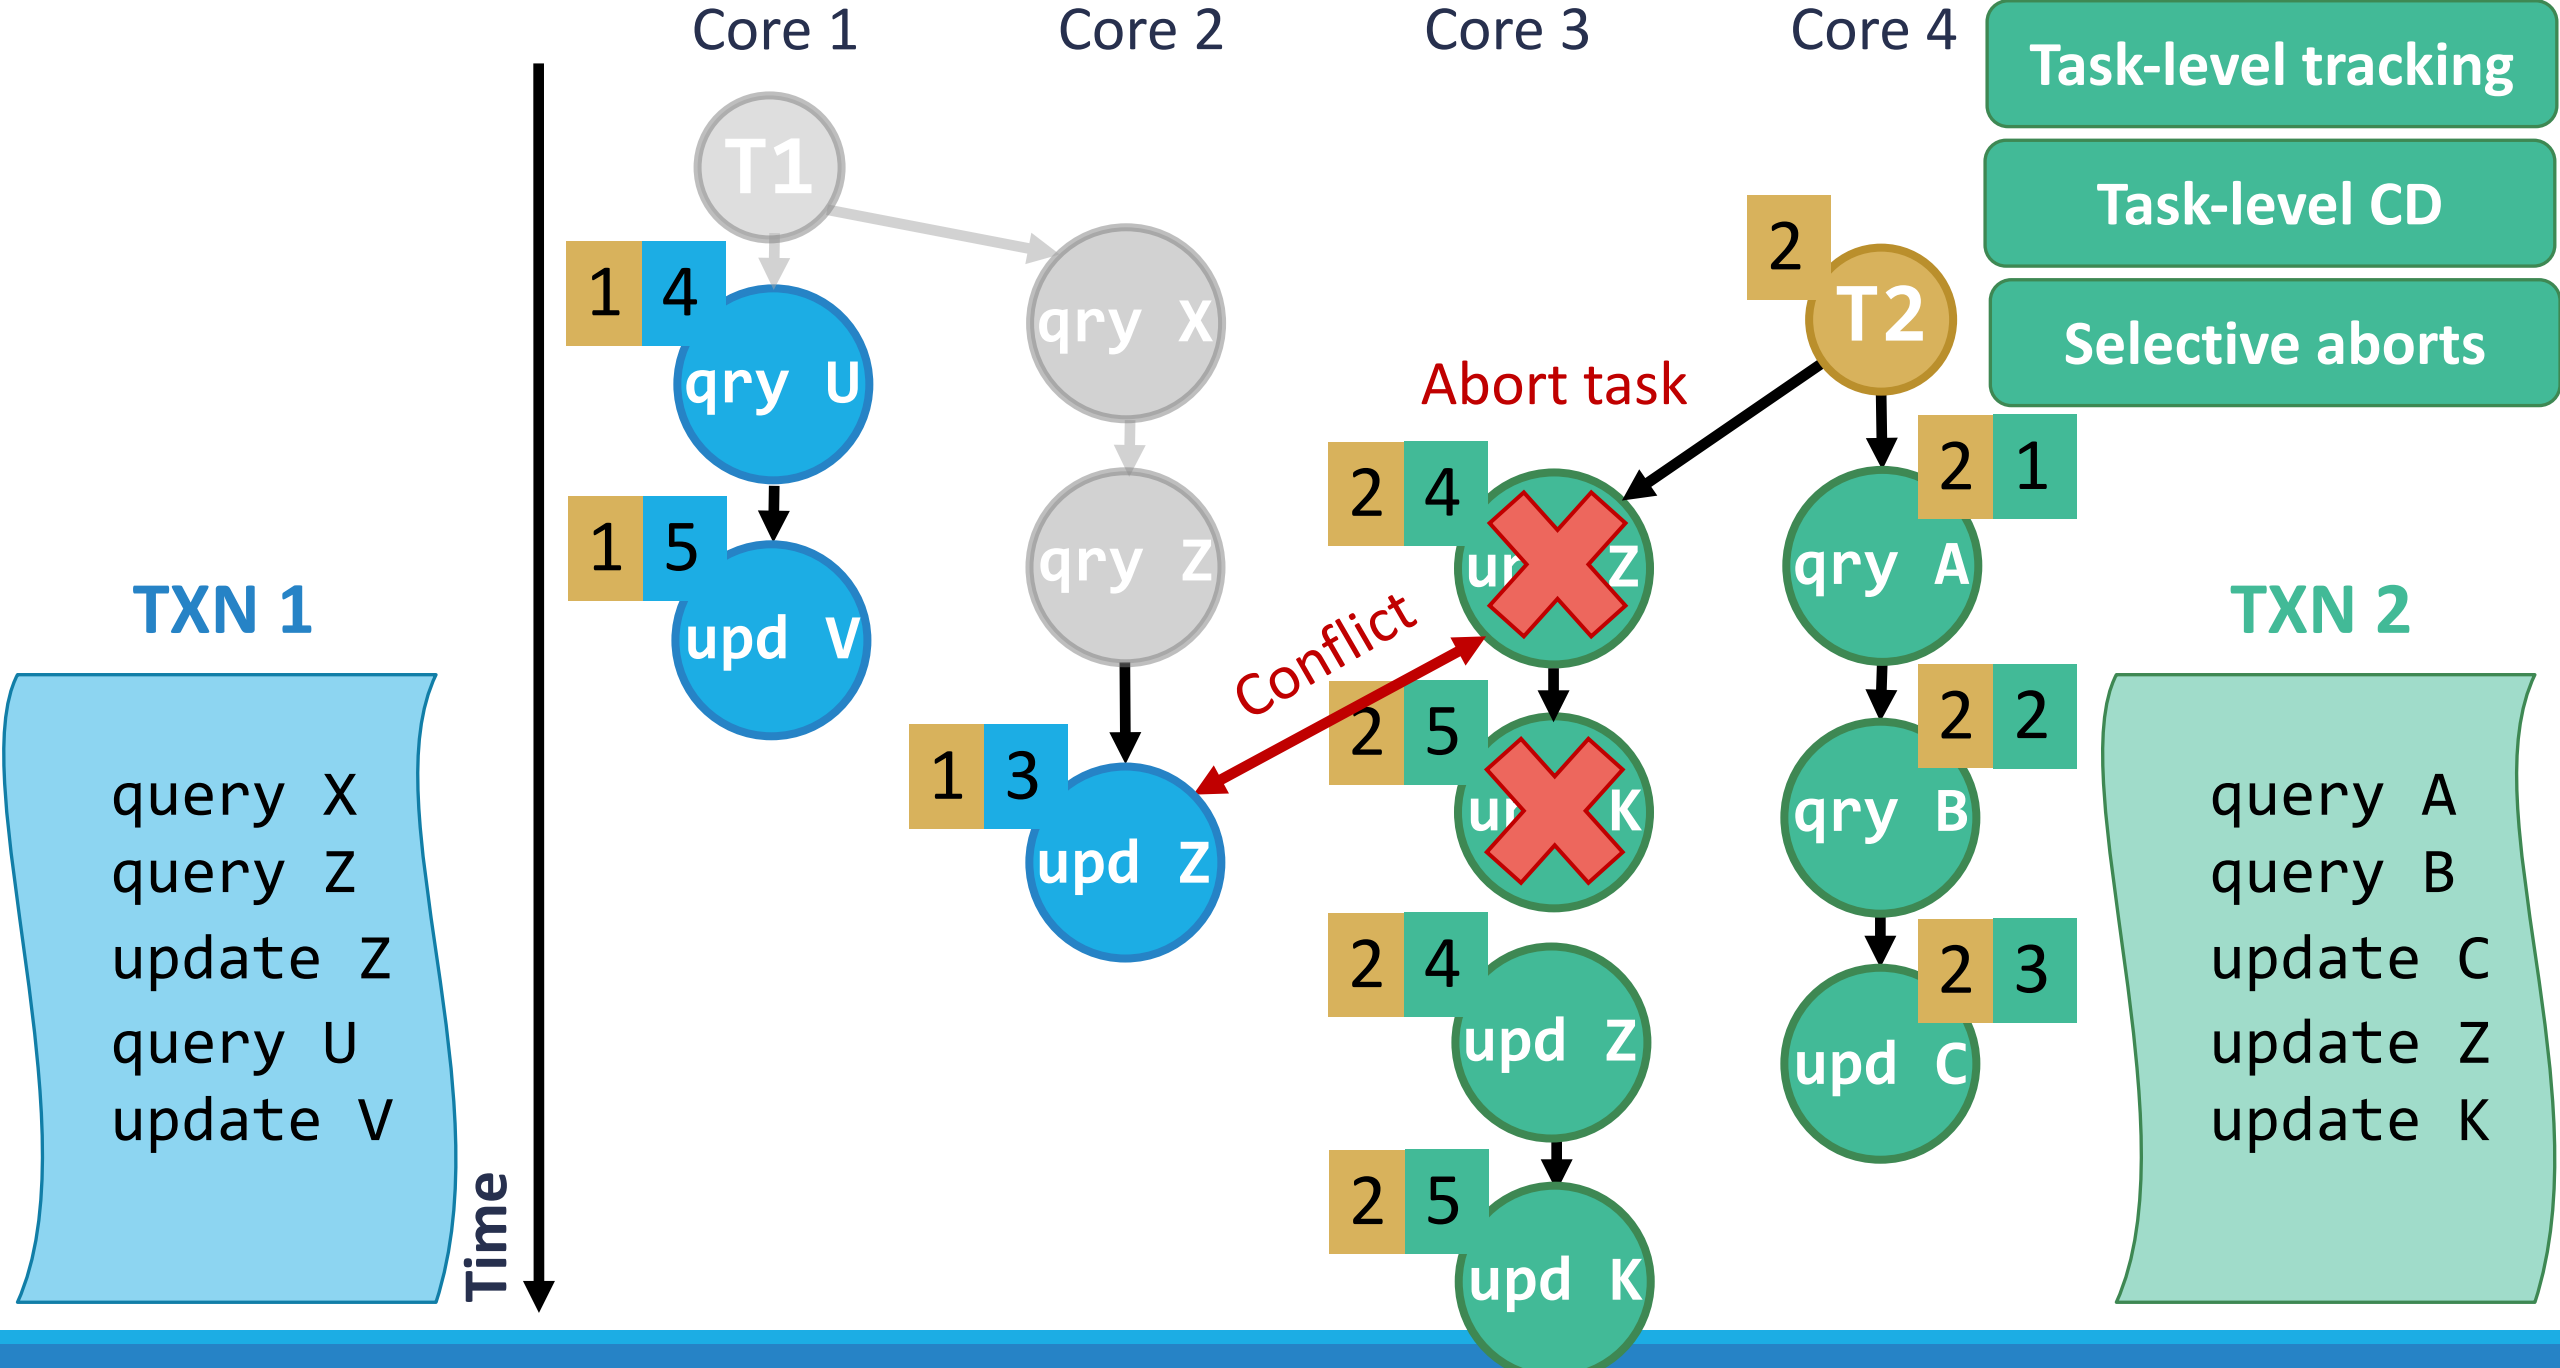


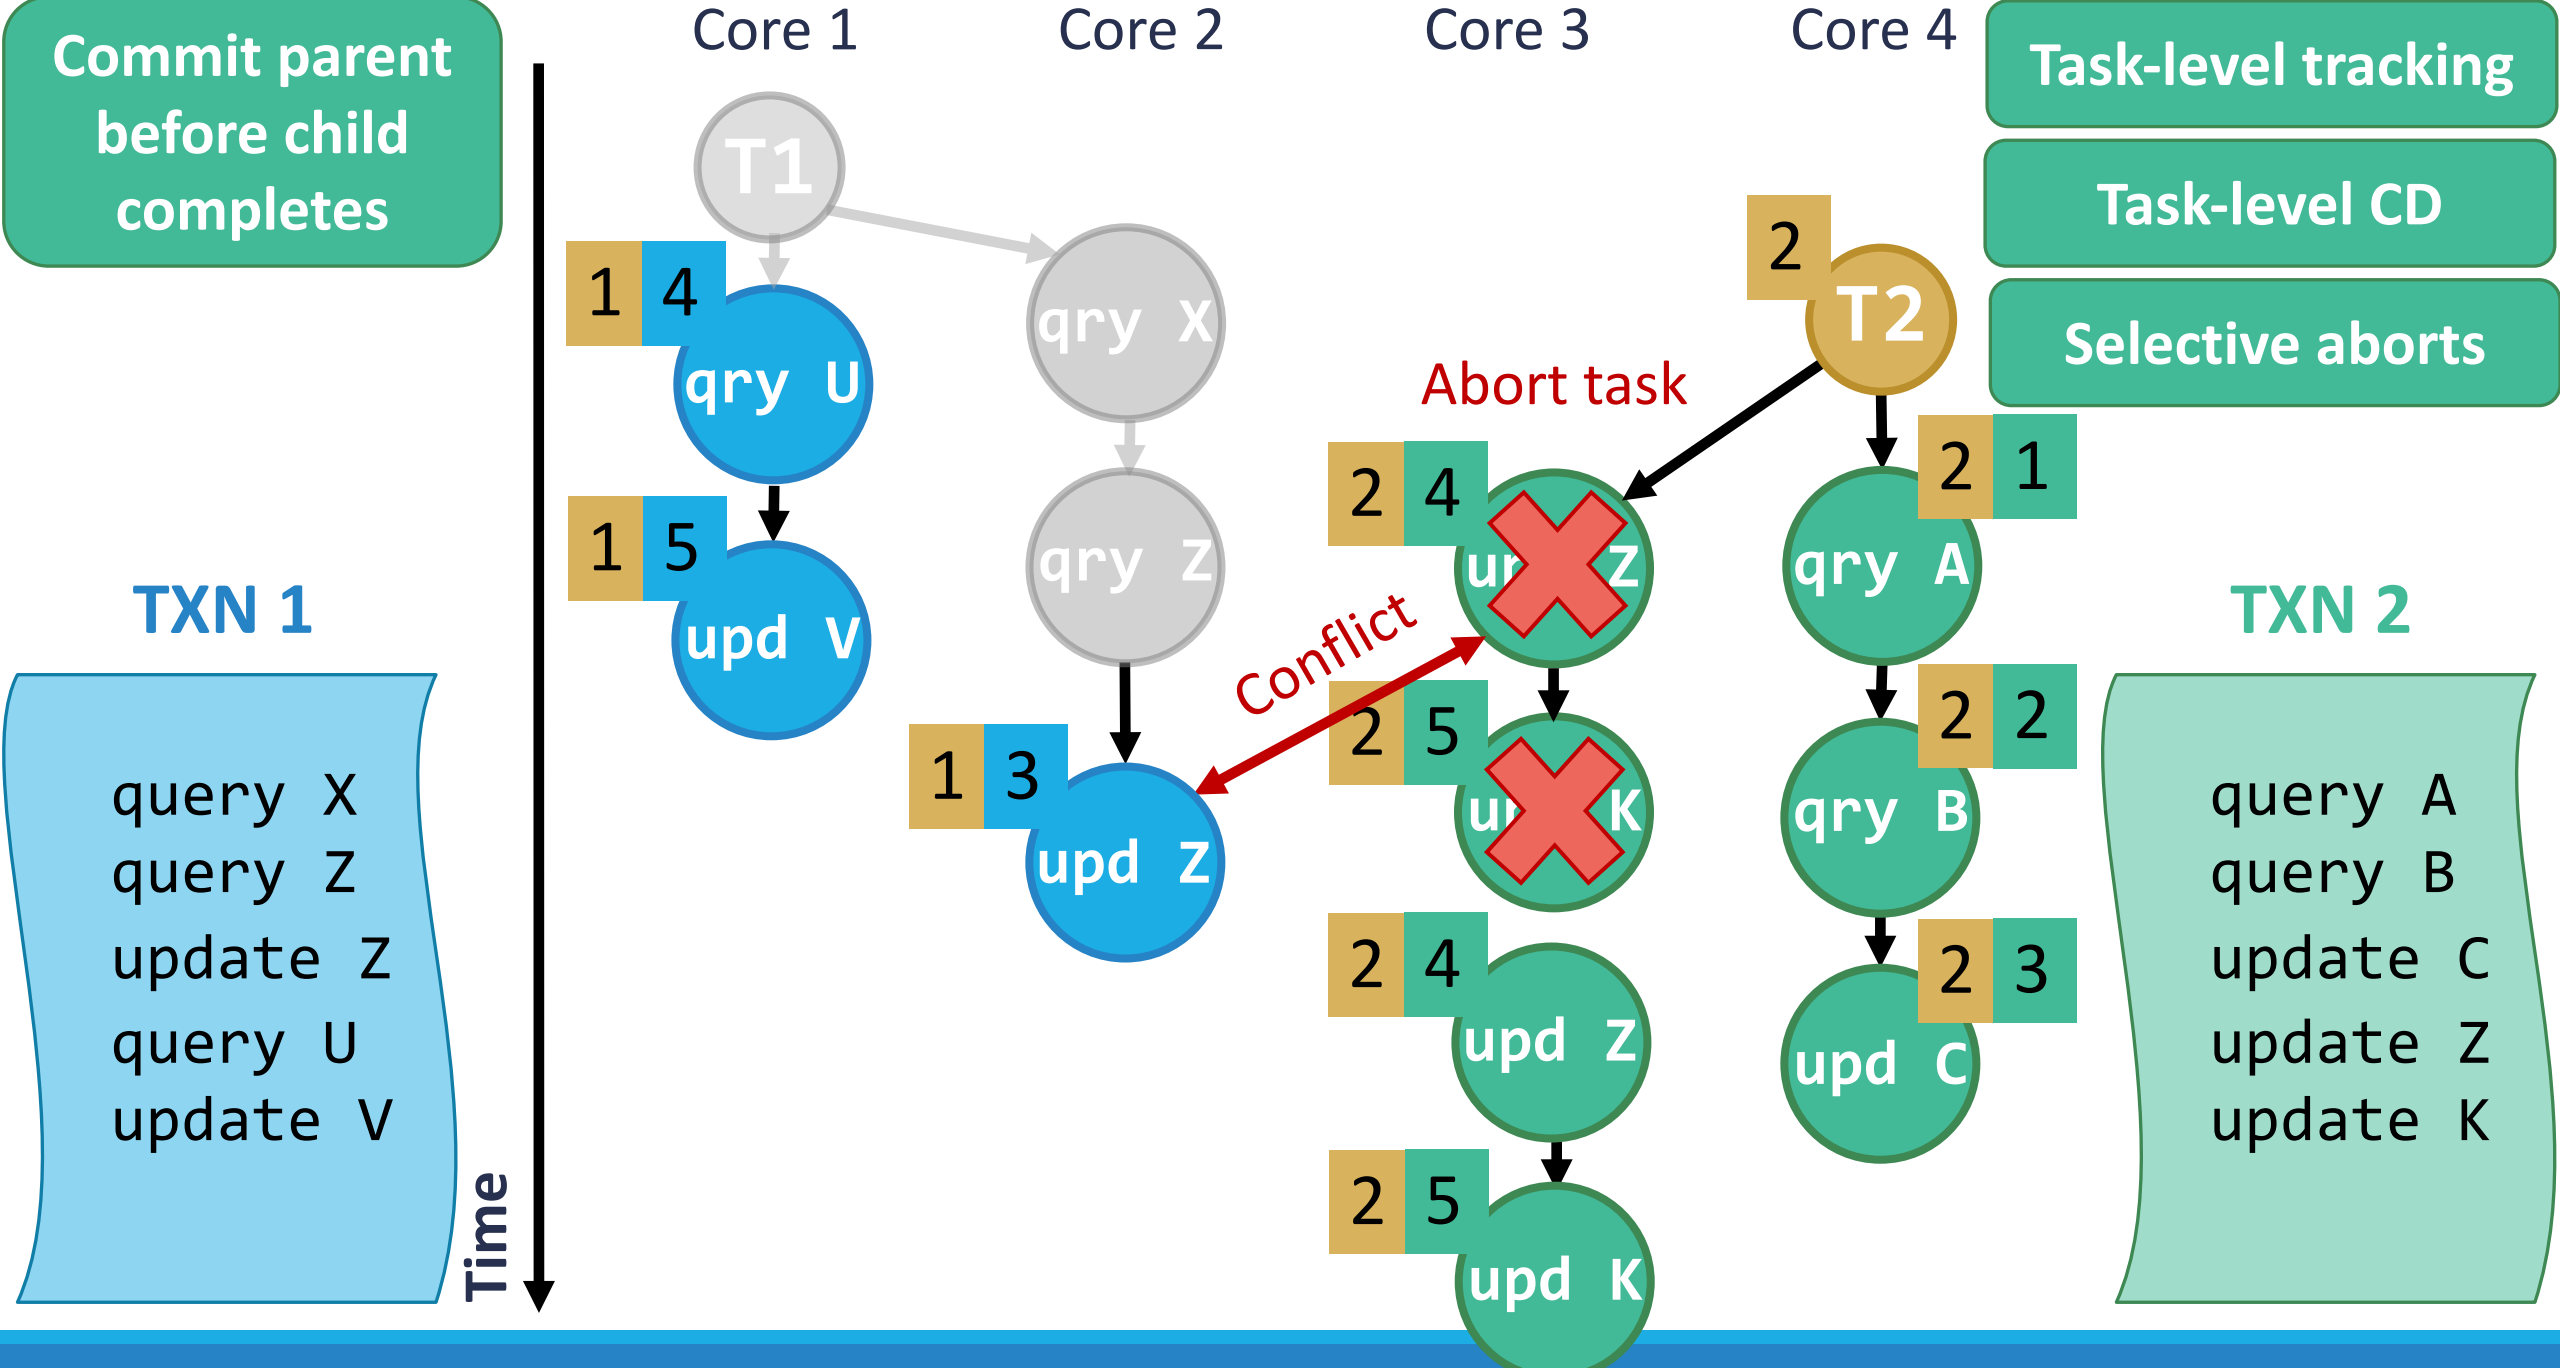












Commit parent  
before child  
completes

Core 1

Core 2

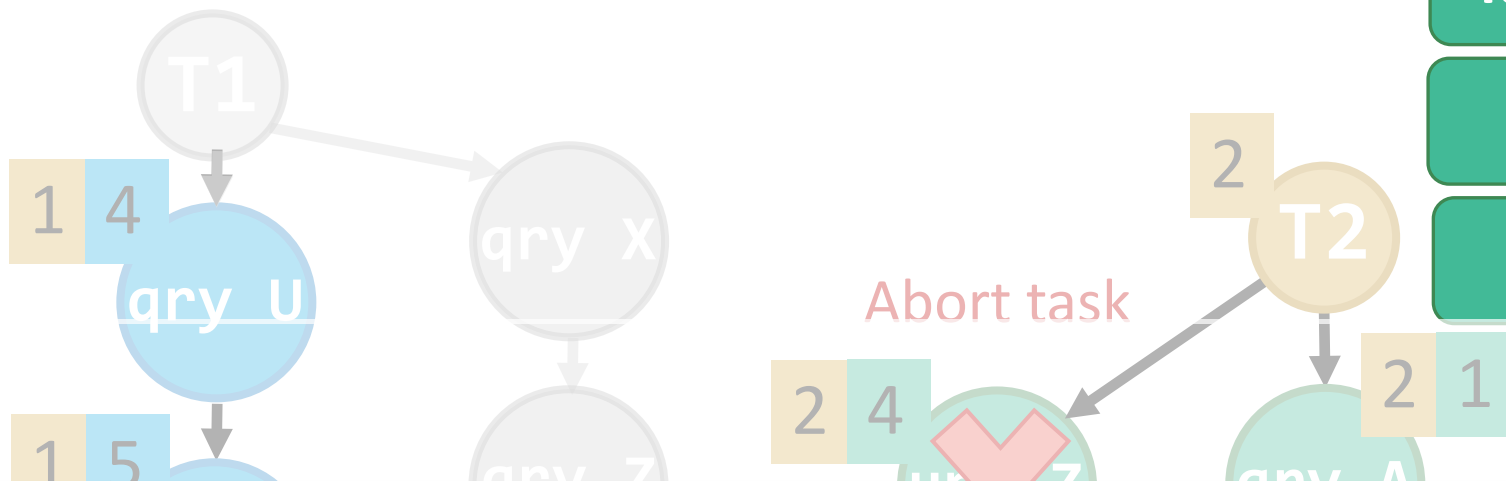
Core 3

Core 4

Task-level tracking

Task-level CD

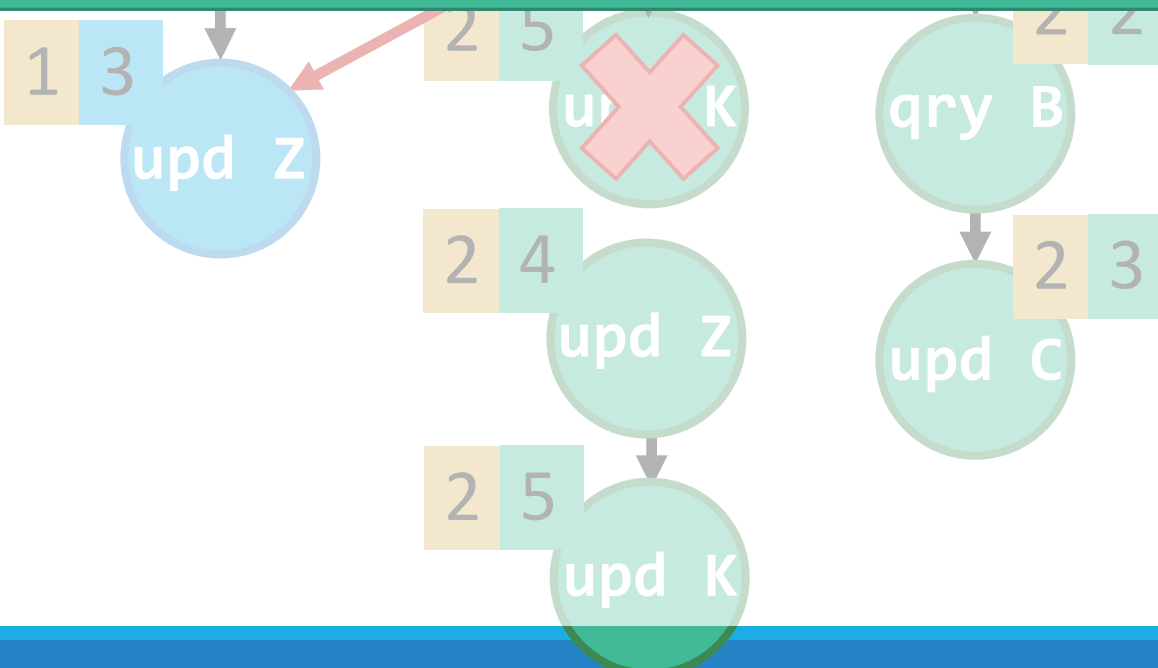
Selective aborts



# Fractal unlocks the benefits of fine-grain parallelism

query X  
query Z  
update Z  
query U  
update V

Time ↓



query A  
query B  
update C  
update Z  
update K



# Evaluation

---

# Methodology

Event-driven, Pin-based simulator

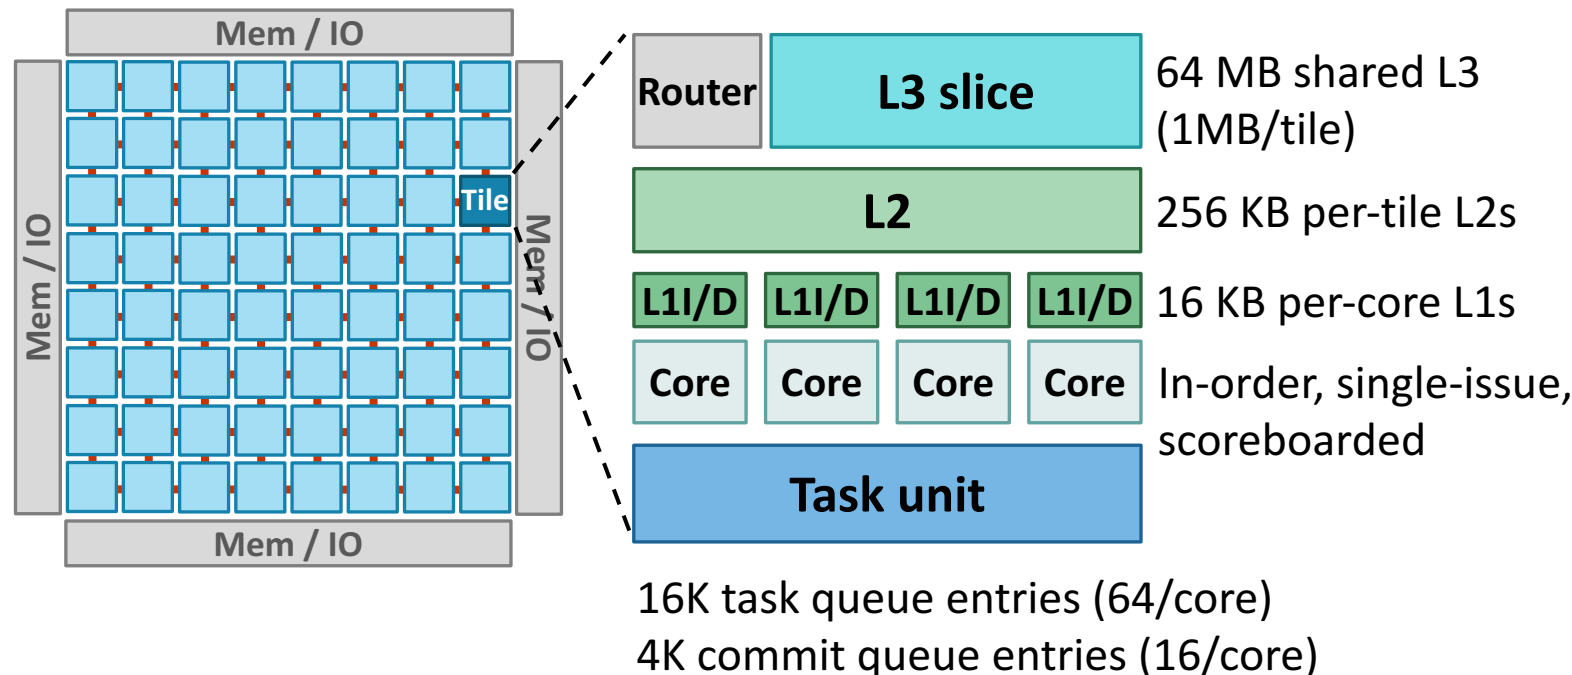
Target system: 256-core, 64-tile chip

Scalability experiments from 1–256 cores

- Scaled-down systems have fewer tiles


## Applications

- Unordered (STAMP): labyrinth, bayes
- Ordered: color, msf, silo, maxflow, mis




# Fractal uncovers abundant nested parallelism

---



# Fractal uncovers abundant nested parallelism

— Flat — Fractal



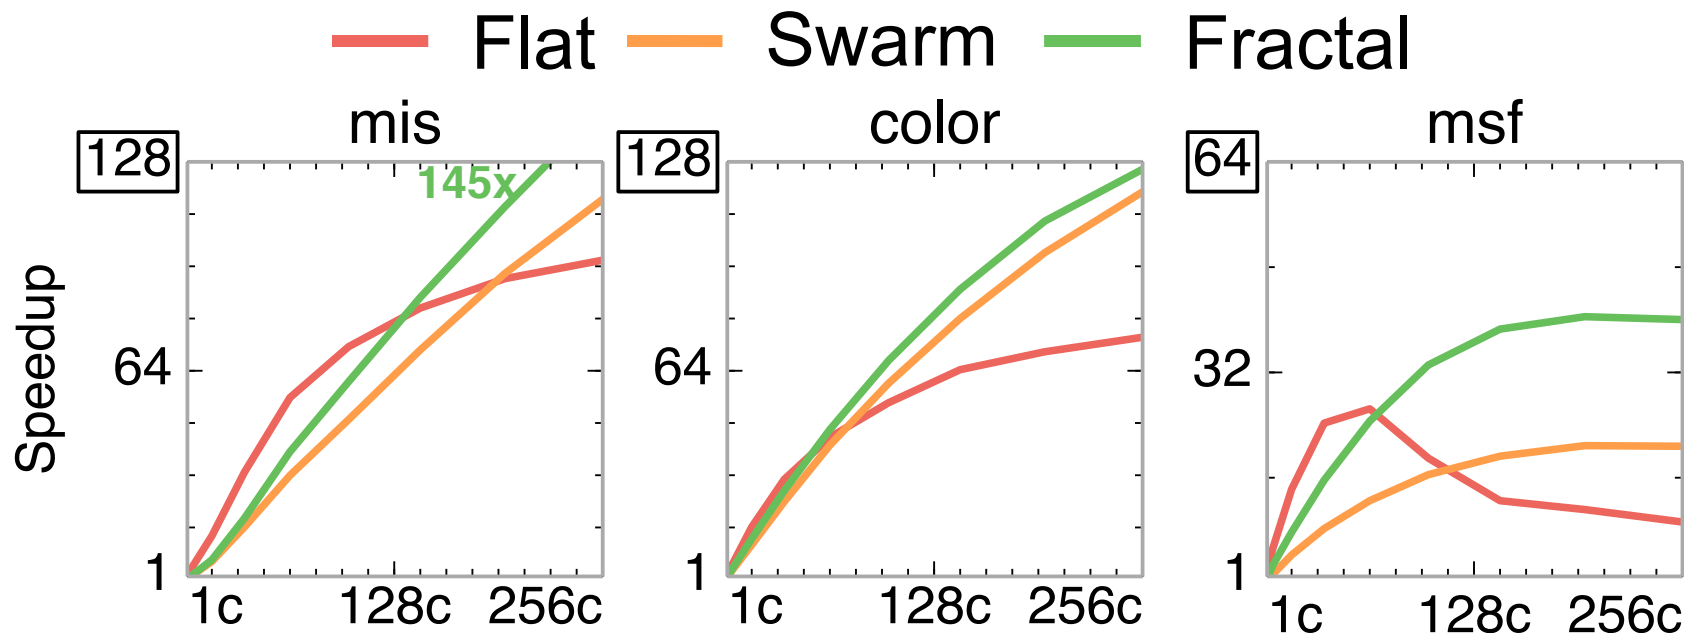
Flat  
1x—4.9x

Fractal  
88x—322x

Flat	3260	1.8 M	16 M
Fractal	373	3590	220

} Average task length  
(cycles)

# Fractal avoids over-serialization



**Flat**  
26x—98x

**Swarm**  
21x—119x

**Fractal**  
40x—145x

<b>Flat</b>	162	633	113
<b>Fractal</b>	115	96	49

} Average task length  
(cycles)

# Conclusion

---

Speculative systems must extract nested parallelism in order to scale large, complex, real-world applications

**Fractal:** An execution model for fine-grain nested speculative parallelism

- Decouple atomicity from parallelism
- Guarantee atomicity by ordering tasks

**Fractal** unlocks the benefits of fine-grain speculative parallelism

- Parallelizes many challenging workloads
- Enables composition of speculative parallel algorithms

# Thank You! Questions?

---

Speculative systems must extract nested parallelism in order to scale large, complex, real-world applications

**Fractal:** An execution model for fine-grain nested speculative parallelism

- Decouple atomicity from parallelism
- Guarantee atomicity by ordering tasks

**Fractal** unlocks the benefits of fine-grain speculative parallelism

- Parallelizes many challenging workloads
- Enables composition of speculative parallel algorithms