

# Locking for Concurrent Transactions on Ontologies

Matthias Thimm

*joint work with Stefan Scheglmann, Steffen Staab, and Gerd Gröner*



### Definition:

#### Transaction

- begin of transaction
- sequence of operations
- end of transaction


#### Critical Operations





- tell
- forget
- ask

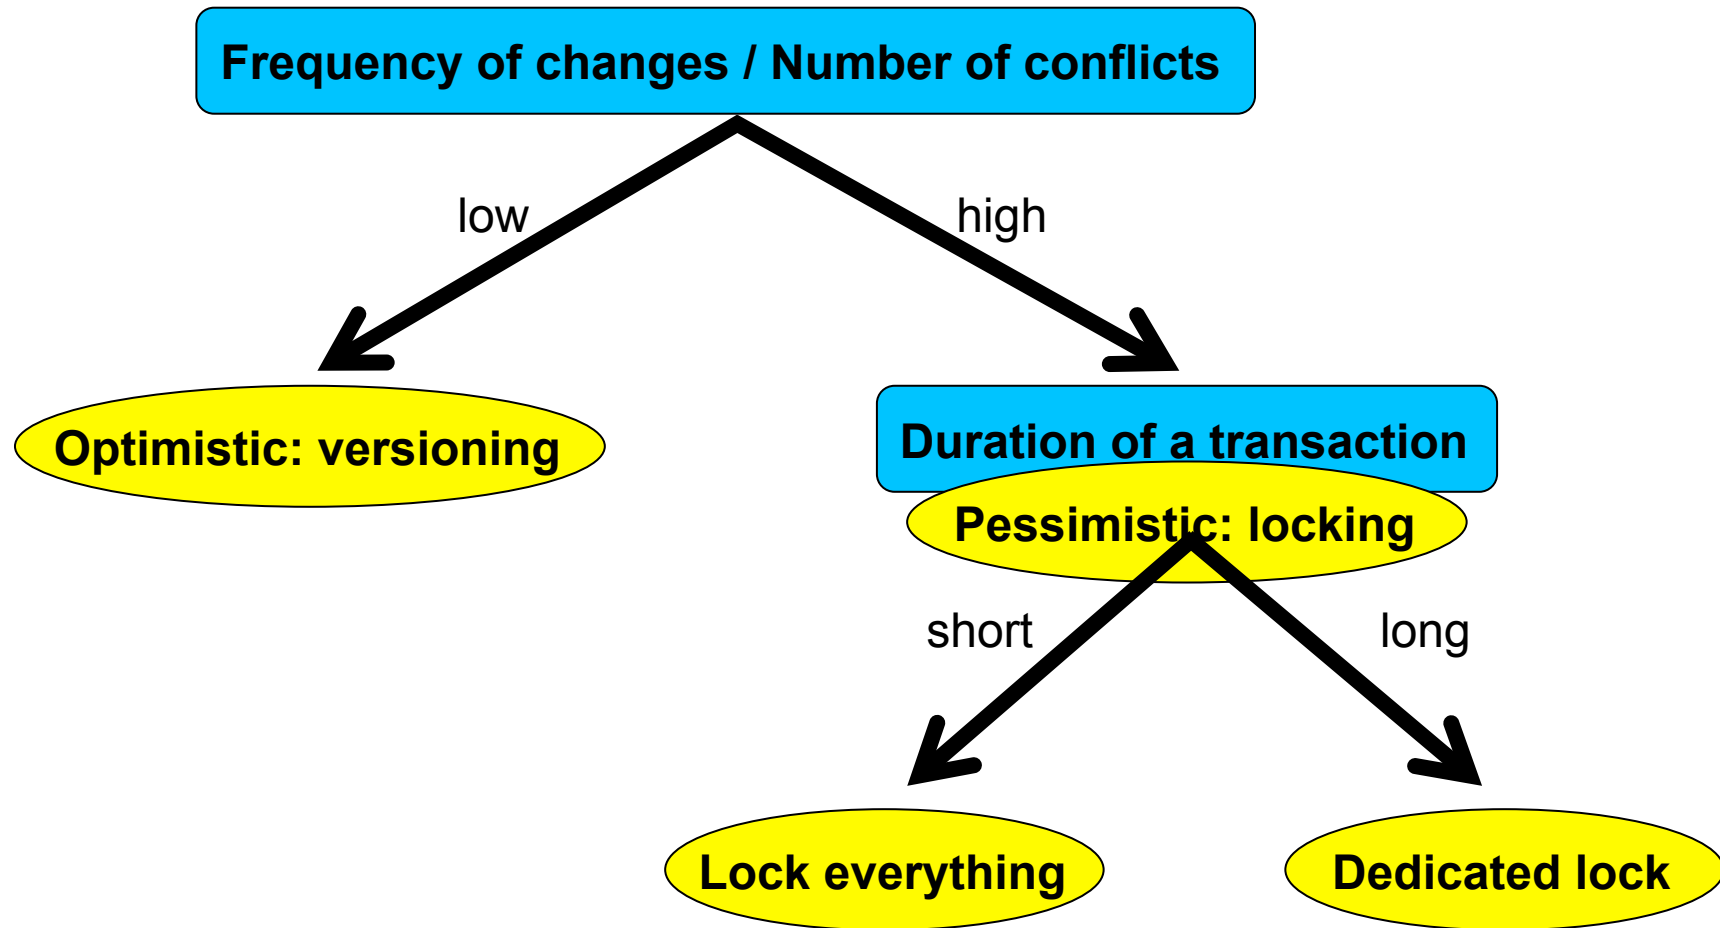
#### Non-critical operation

- user input
- response
- etc

critical operation











forget( $B \sqsubseteq D_1$ ), tell( $B \sqsubseteq D_2$ )



$$A_1 \equiv \forall R.D_1$$

$$A_2 \equiv \forall R.D_2$$

$$A \sqsubseteq \forall R.B$$

$$D_1 \sqcap D_2 \equiv \perp$$

$$B \sqsubseteq D_2$$




ask( $A \sqsubseteq ?X$ )



$A \sqsubseteq A_2$

Alf Before

Bob

$$A \sqsubseteq A_2 \sqsubseteq T$$



$$A_1 \equiv \forall R.D_1$$

$$A_2 \equiv \forall R.D_2$$

$$A \sqsubseteq \forall R.B$$

$$D_1 \sqcap D_2 \equiv \perp$$

$$B \sqsubseteq D_2$$

Alf	Before	Interleaved	Bob
	$A \sqsubseteq A_2 \sqsubseteq T$	$A \sqsubseteq T$	

After

WeST 



update!

Alf

forget( $B \sqsubseteq D_1$ ), tell( $B \sqsubseteq D_2$ )





ask!

Bob

ask( $A \sqsubseteq ?X$ )

$A \sqsubseteq A_1$

$A_1 \equiv \forall R.D_1$   
 $A_2 \equiv \forall R.D_2$   
 $A \sqsubseteq \forall R.B$   
 $D_1 \sqcap D_2 \equiv \perp$   
 $B \sqsubseteq D_2$

Alf	Before	Interleaved	After	Bob
-----	--------	-------------	-------	-----

$A \sqsubseteq A_2 \sqsubseteq T$

$A \sqsubseteq T$



$A \sqsubseteq A_1 \sqsubseteq T$

# LOCKING FOR ONTOLOGY-ACCESS





forget( $B \sqsubseteq D_1$ )  
begin



$A_1 \equiv \forall R.D_1$   
 $A_2 \equiv \forall R.D_2$   
 $A \sqsubseteq \forall R.B$   
 $D_1 \sqcap D_2 \equiv \perp$   
 $B \sqsubseteq D_1$   
 $X \sqsubseteq Y$

## Core idea: Ontology module based locks

- Syntactical locality based module approximation
- Function  $S()$ 
  - given a set of axioms
  - provides us with a signature
- Function  $M()$ 
  - given a signature
  - provides us with a module for  $S$




update!


Alf

tell( $B \sqsubseteq D_2$ )  
forget( $B \sqsubseteq D_1$ )




$$M(S(\{B \sqsubseteq D_1\})) \cap M(S(\{A \sqsubseteq ?X\})) \neq \emptyset$$





- $A_1 \equiv \forall R.D_1$
- $A_2 \equiv \forall R.D_2$
- $A \sqsubseteq \forall R.B$
- $D_1 \sqcap D_2 \equiv \perp$
- $B \sqsubseteq D_2$
- $X \sqsubseteq Y$



ask!  
 $A \sqsubseteq A_2$

Bob

ask( $A \sqsubseteq ?X$ )  
 $A \sqsubseteq A_2$

$$M(S(A \sqsubseteq ?X))$$

## Input:

- National Cancer Thesaurus (NCIt)
  - OWL EL++ Ontology
  - 4 consecutive versions
  - ~600K Axioms,
    - 500K annotation, 38K classes, 90 object properties

## Steps:


1. Identifying 420 Transactions of 6-12 operations
  - How: From syntactic diffs of consecutive versions
2. Each critical operation followed by non-critical with same duration
3. Constructing Histories from 2-4 Transactions

## Output:

1200 Histories consisting of 24 – 96 atomic operations each



- 30% (~240) of the histories are serializable
  - Due to random generation of transactions
- Average serializable history 76,6% of the steps of serial execution
- In average 2.832 sec for a single lock calculation
- Linux Virtual Machine 8Gig Ram, 1 Core  
(Dual Xeon Hexacore, 2.9Ghz, 96 Gig Ram running ~20 VMs)



- 10% of module calculation time for the module of the upcoming operation
- 90% to recalculate the lock

**Incremental module calculation methods would lead to a downscale**

- Incremental module calculation
- Pre-computation strategies
- Real-world evaluation
- Integration into ontology development tool

*Thank you for your attention.*

See also [Stefan Scheglmann, Steffen Staab, Matthias Thimm, Gerd Gröner. Locking for Concurrent Transactions on Ontologies. In Proceedings of the 10th Extended Semantic Web Conference (ESWC'13). Montpellier, May, 2013.]