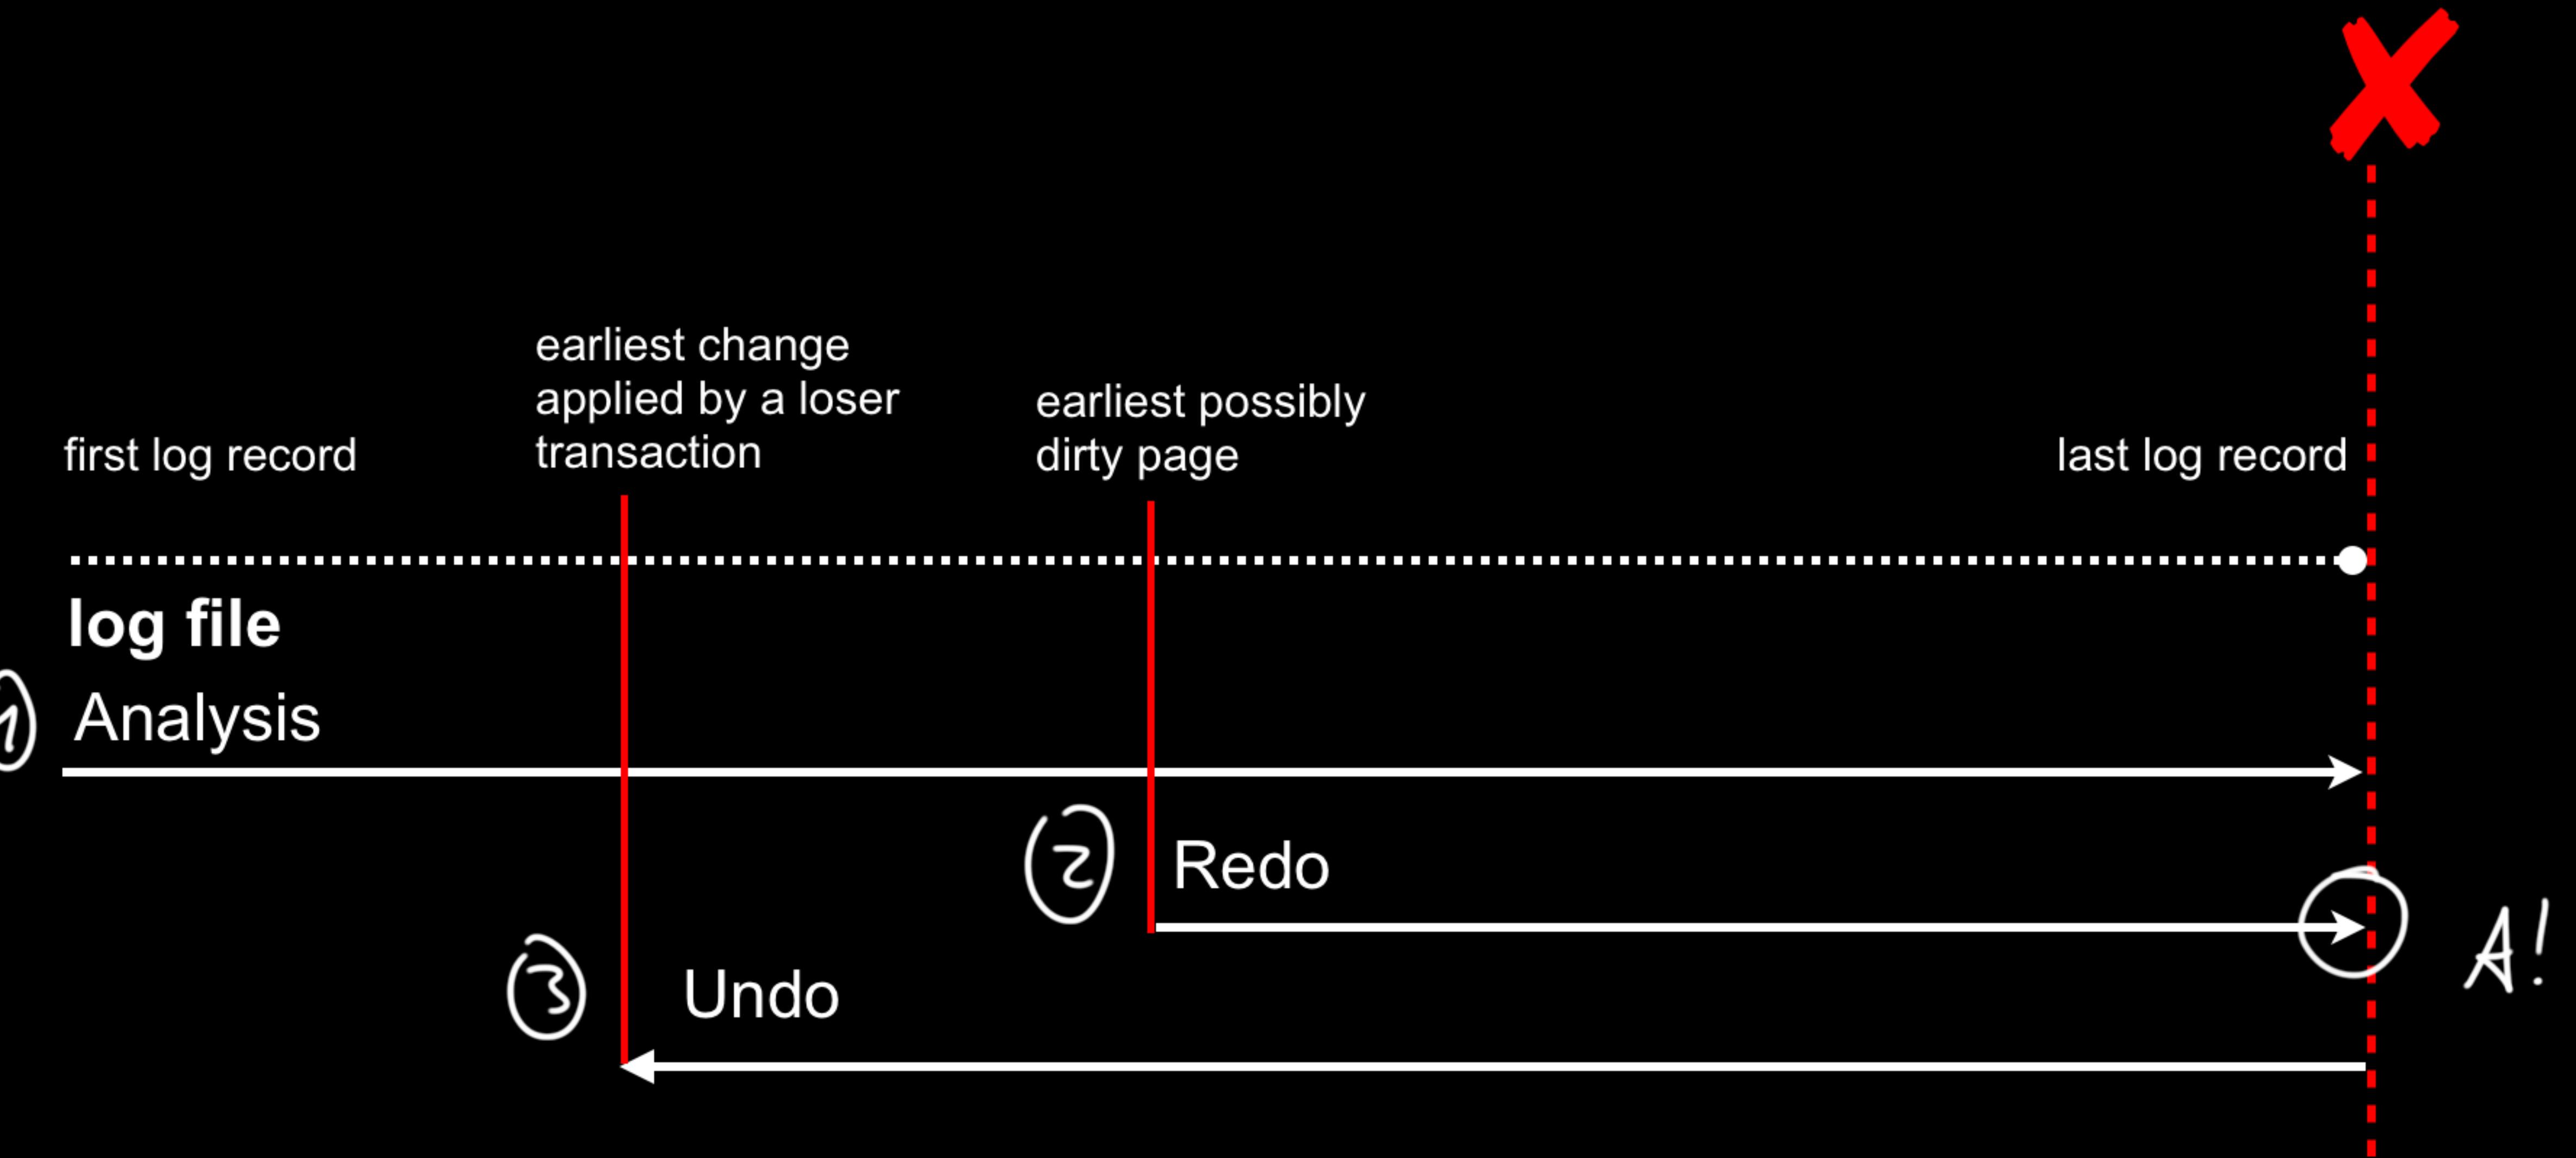


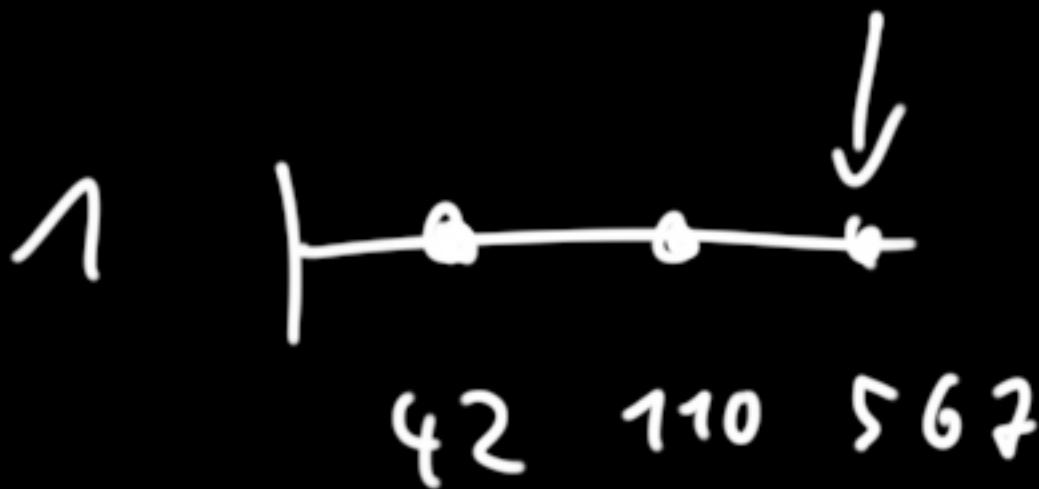
# 3 Phases of ARIES



notice: for the moment no checkpointing!

# Transaction Table TT

TT	
TaID	lastLSN
1	567
2	42
7	67
3	12



# Dirty Page Table DPT

DPT	
pageID	recoveryLSN
42	567
46	568
77	34
3	42

TaID: key of the transaction

lastLSN: LSN of the most recent log record seen for this transaction, i.e. the **latest** change done by this transaction

# Transaction Table TT

TT	
TaID	lastLSN
1	567
2	42
7	67
3	12

TaID: key of the transaction

lastLSN: LSN of the most recent log record seen for this transaction, i.e. the **latest** change done by this transaction

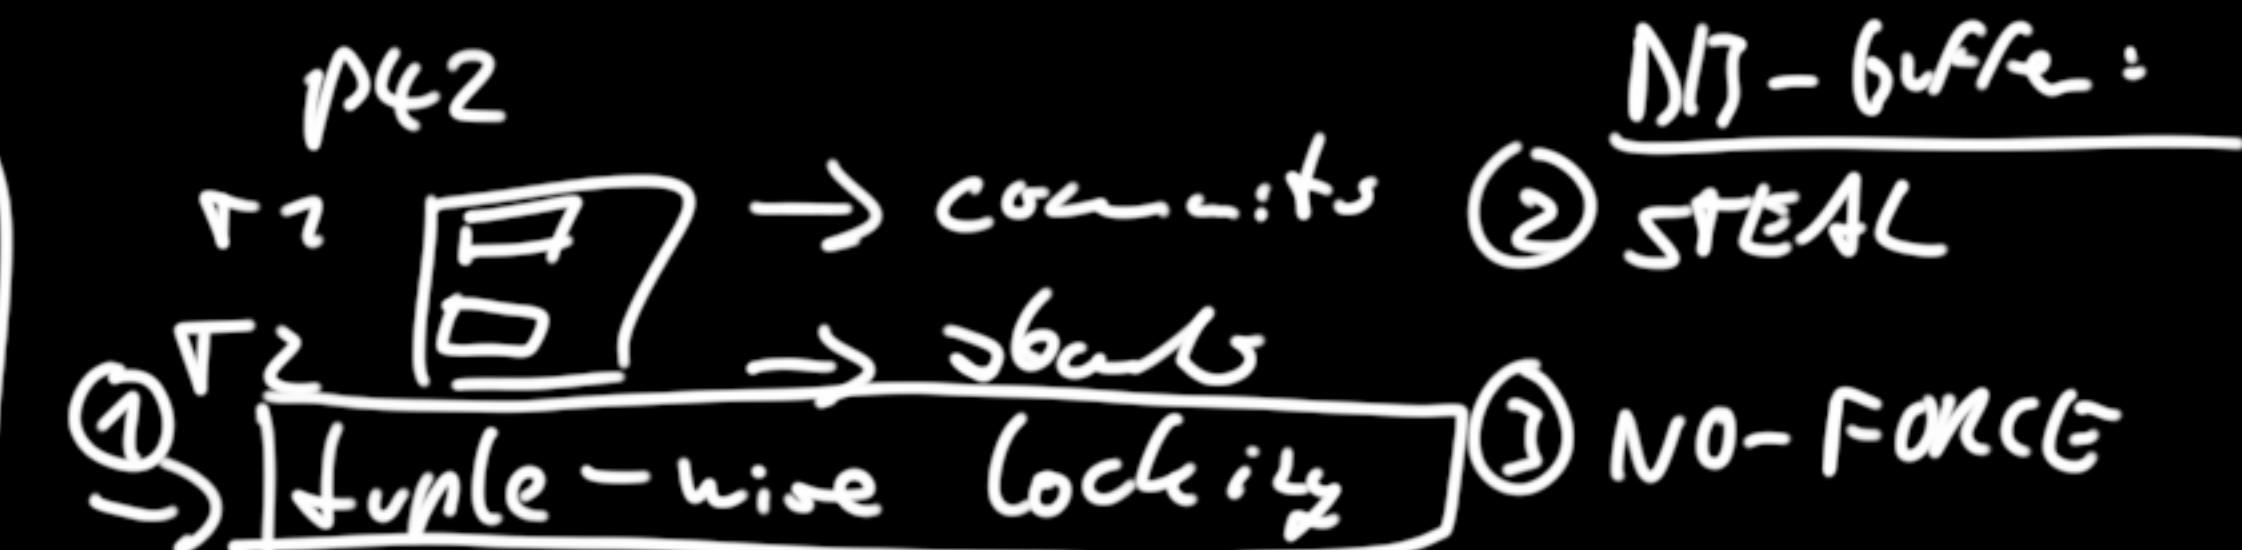
**assumptions**: page may contain changes from losers and winners, i.e. two concurrent transactions may operate **on the same page**

# Dirty Page Table DPT

DPT	
pageID	recoveryLSN
42	567
46	568
77	34
3	42

pageID: key of a page

recoveryLSN: LSN of **first** log record that made this page dirty, i.e. the **earliest** change done to this page



# The Log Records in Detail

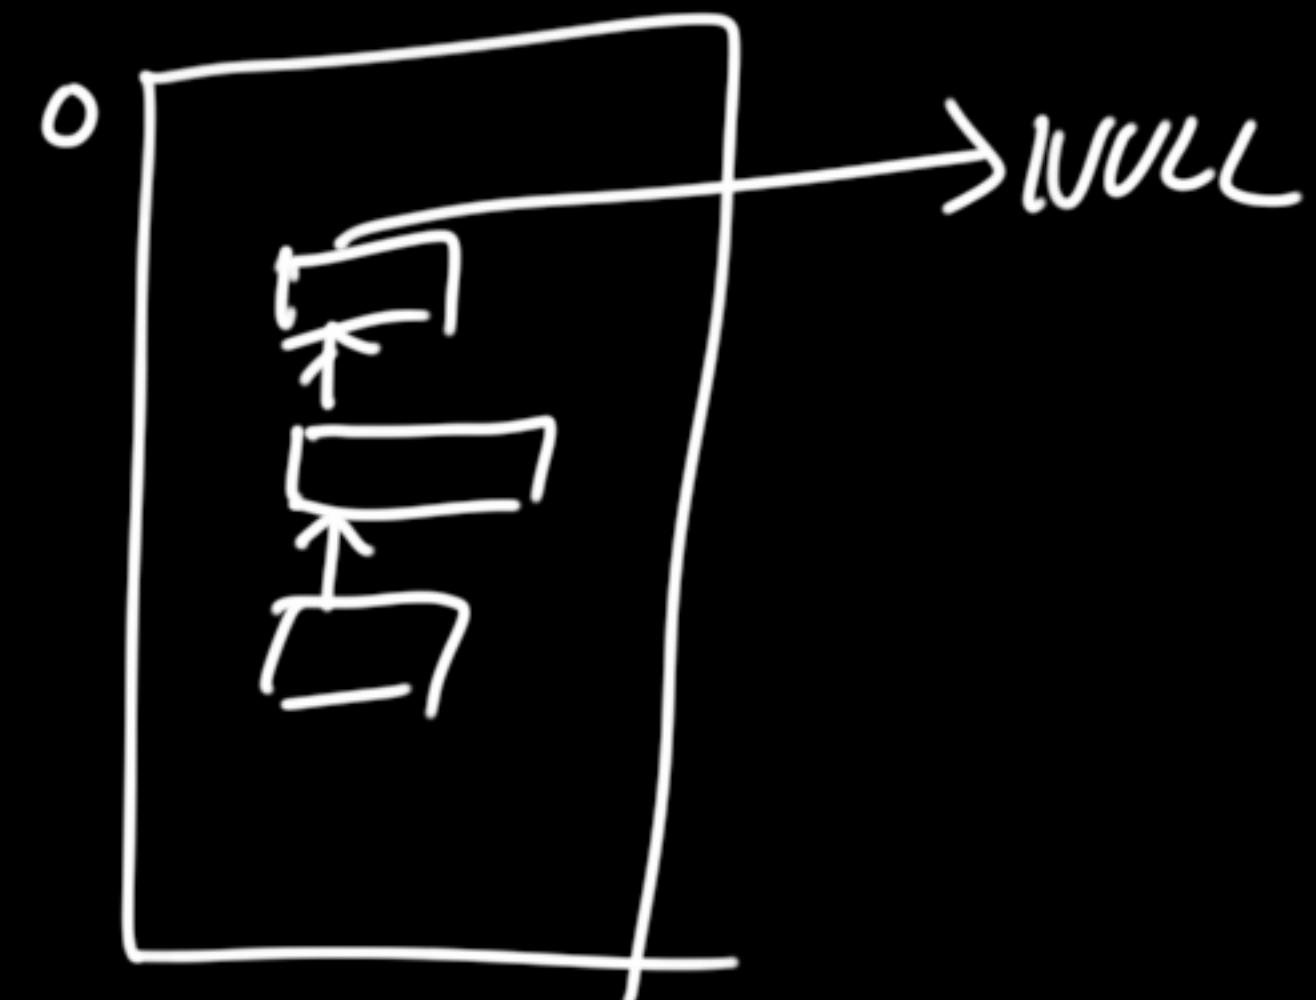
All Log Records:

[prevLSN, TaID, type]

420 | 42 |

offset = LSN

562



prevLSN: (Previous Log Sequence Number): link to the previous LSN for this TA (if not first entry)

TaID: foreign key to unique transaction ID

# The Log Records in Detail

All Log Records:

[prevLSN, TaID, type]

prevLSN: (Previous Log Sequence Number): link to the previous LSN for this TA (if not first entry)

TaID: foreign key to unique transaction ID

type: kind of log record; either: update, commit or compensation (for undoing updates)

# The Log Records in Detail: Update Log Records

All Log Records:

[prevLSN, Taid, type]

Update Log Record:

[prevLSN, Taid, “update”, pageID, redo info, undo info]

WAL?

# The Log Records in Detail: Update Log Records

All Log Records:

[prevLSN, Taid, type]

Update Log Record:

[prevLSN, Taid, “update”, pageID, redo info, undo info]

physical      or      logical

pageID: foreign key to pageID

# The Log Records in Detail: Update Log Records

All Log Records:

[prevLSN, Taid, type]

Update Log Record:

[prevLSN, Taid, “update“, pagID, redo info, undo info]

pagID: foreign key to pagID

redo info: information how to redo the change reflected by this log record (physiological)

# The Log Records in Detail: Update Log Records

All Log Records:

Logical undo : concurrency in  
index structure

[prevLSN, TaID, type]

Update Log Record:

[prevLSN, TaID, “update”, pageID, redo info, undo info]

pageID: foreign key to pagID

for the moment ↗

ARIES : redo : physiological  
undo : (logical) physiological

redo info: information how to redo the change reflected by this log record (physiological)

undo info: information how to undo the change reflected by this log record (logical)

# Example Log File

LSN ~ version

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info]

“update“

DB buffer

page 42

LSN=-
a=77
b=55

page 46

LSN=-
c=22

log



TT	
TaID	lastLSN

DPT	
pageID	recoveryLSN

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info]

“update“

DB buffer

page 42

LSN=1
a=78
b=55

page 46

LSN=-
c=22

log

1: [-, 1, “update“, 42, a+=1, a-=1]

TT	
TaID	lastLSN
1	1

DPT	
pageID	recoveryLSN
42	1

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pagelD, redo info, undo info]

“update“

DB buffer

page 42

LSN=2
a=78
b=58

page 46

LSN=-
c=22

log

- 1: [-, 1, “update“, 42, a+=1, a-=1]
- 2: [-, 2, “update“, 42, b+=3, b-=3]

TT	
TaID	lastLSN
1	1
2	2

DPT	
pagelD	recoveryLSN
42	1

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info]

“update“

DB buffer

page 42

LSN=2
a=78
b=58

page 46

LSN=3
c=24

log

- 1: [-, 1, “update“, 42, a+=1, a-=1]
- 2: [-, 2, “update“, 42, b+=3, b-=3]
- 3: [2, 2, “update“, 46, c+=2, c-=2]

TT	
TaID	lastLSN
1	1
2	3

DPT	
pageID	recoveryLSN
42	1
46	3

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info]

“update“

DB buffer

page 42

LSN=4
a=78
b=59

page 46

LSN=3
c=24

log

- 1: [-, 1, “update“, 42, a+=1, a-=1]
- 2: [-, 2, “update“, 42, b+=3, b-=3]
- 3: [2, 2, “update“, 46, c+=2, c-=2]
- 4: [1, 1, “update“, 42, b+=1, b-=1]

TT	
TaID	lastLSN
1	4
2	3

DPT	
pageID	recoveryLSN
42	1
46	3

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info]

“update“

DB buffer

page 42

LSN=4
a=78
b=59

page 46

LSN=3
c=24

log

- 1: [-, 1, “update“, 42, a+=1, a-=1]
- 2: [-, 2, “update“, 42, b+=3, b-=3]
- 3: [2, 2, “update“, 46, c+=2, c-=2]
- 4: [1, 1, “update“, 42, b+=1, b-=1]
- 5: [3, 2, “commit“]

TT	
TaID	lastLSN
1	4

DPT	
pageID	recoveryLSN
42	1
46	3

# The Log Records in Detail: Compensation Log Records

All Log Records:

[prevLSN, Taid, type]

Update Log Record:

[prevLSN, Taid, “update“, pagelD, redo info, undo info]

Compensation Log Record (CLR, for undoing updates):

[prevLSN, Taid, “compensation“, **redoTheUndo** info, undoNextLSN]

**redoTheUndo** info: information how to redo the undo (reflected by the log record creating this CLR)

① Analyse  
② Redo → handle CLR  
③ Undo for every undo operation:  
write CLR

# The Log Records in Detail: Compensation Log Records

All Log Records:

[prevLSN, Taid, type]

Update Log Record:

[prevLSN, Taid, “update“, pagID, redo info, undo info]

Compensation Log Record (CLR, for undoing updates):

[prevLSN, Taid, “compensation“, pagID, **redoTheUndo** info, undoNextLSN]

**redoTheUndo** info: information how to redo the undo (reflected by the log record creating this CLR)

**undoNextLSN**: link to the next log record to be undone for this transaction

i.e., if a log record LR wrote this CLR, we set CLR.undoNextLSN := LR.prevLSN

In general, and as mentioned above, in the multi-level ARIES method, CLRs may be (fully) logical not showing a pagID.

However, I simplified the explanation of ARIES in this video such that every log record is either physical or physiological, but NOT logical.

Therefore, the schema of the CLRs show a pagID (as it does in the following example).

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, "update", pageID, redo info, undo info]

"update"

LSN: [prevLSN, TaID, "compensation", redoTheUndo info, undoNextLSN]

"compensation"

DB buffer

page 42

LSN=4
a=78
b=59

log

- 1: [-, 1, "update", 42, a+=1, a-=1]
- 2: [-, 2, "update", 42, b+=3, b-=3]
- 3: [2, 2, "update", 46, c+=2, c-=2]
- 4: [1, 1, "update", 42, b+=1, b-=1]
- 5: [3, 2, "commit"]

① Analysis

② Roll



TT	
TaID	lastLSN
1	4

page 46

LSN=3
c=24

DPT	
pageID	recoveryLSN
42	1
46	3

max

min LSN

# Example Log File

LSN: [prevLSN, TaID, type]

all

LSN: [prevLSN, TaID, "update", pageID, redo info, undo info]

"update"

LSN: [prevLSN, TaID, "compensation", redoTheUndo info, undoNextLSN]

"compensation"

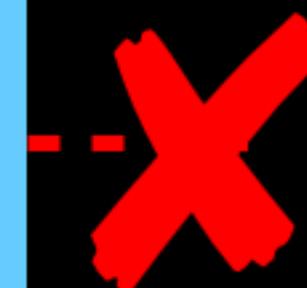
DB buffer

page 42

LSN=6
a=78
b=58

log

- 1: [-, 1, "update", 42, a+=1, a-=1]
- 2: [-, 2, "update", 42, b+=3, b-=3]
- 3: [2, 2, "update", 46, c+=2, c-=2]
- 4: [1, 1, "update", 42, b+=1, b-=1]
- 5: [3, 2, "commit"]
- 6: [4, 1, "compensation", 42, b-=1, 1]



page 46

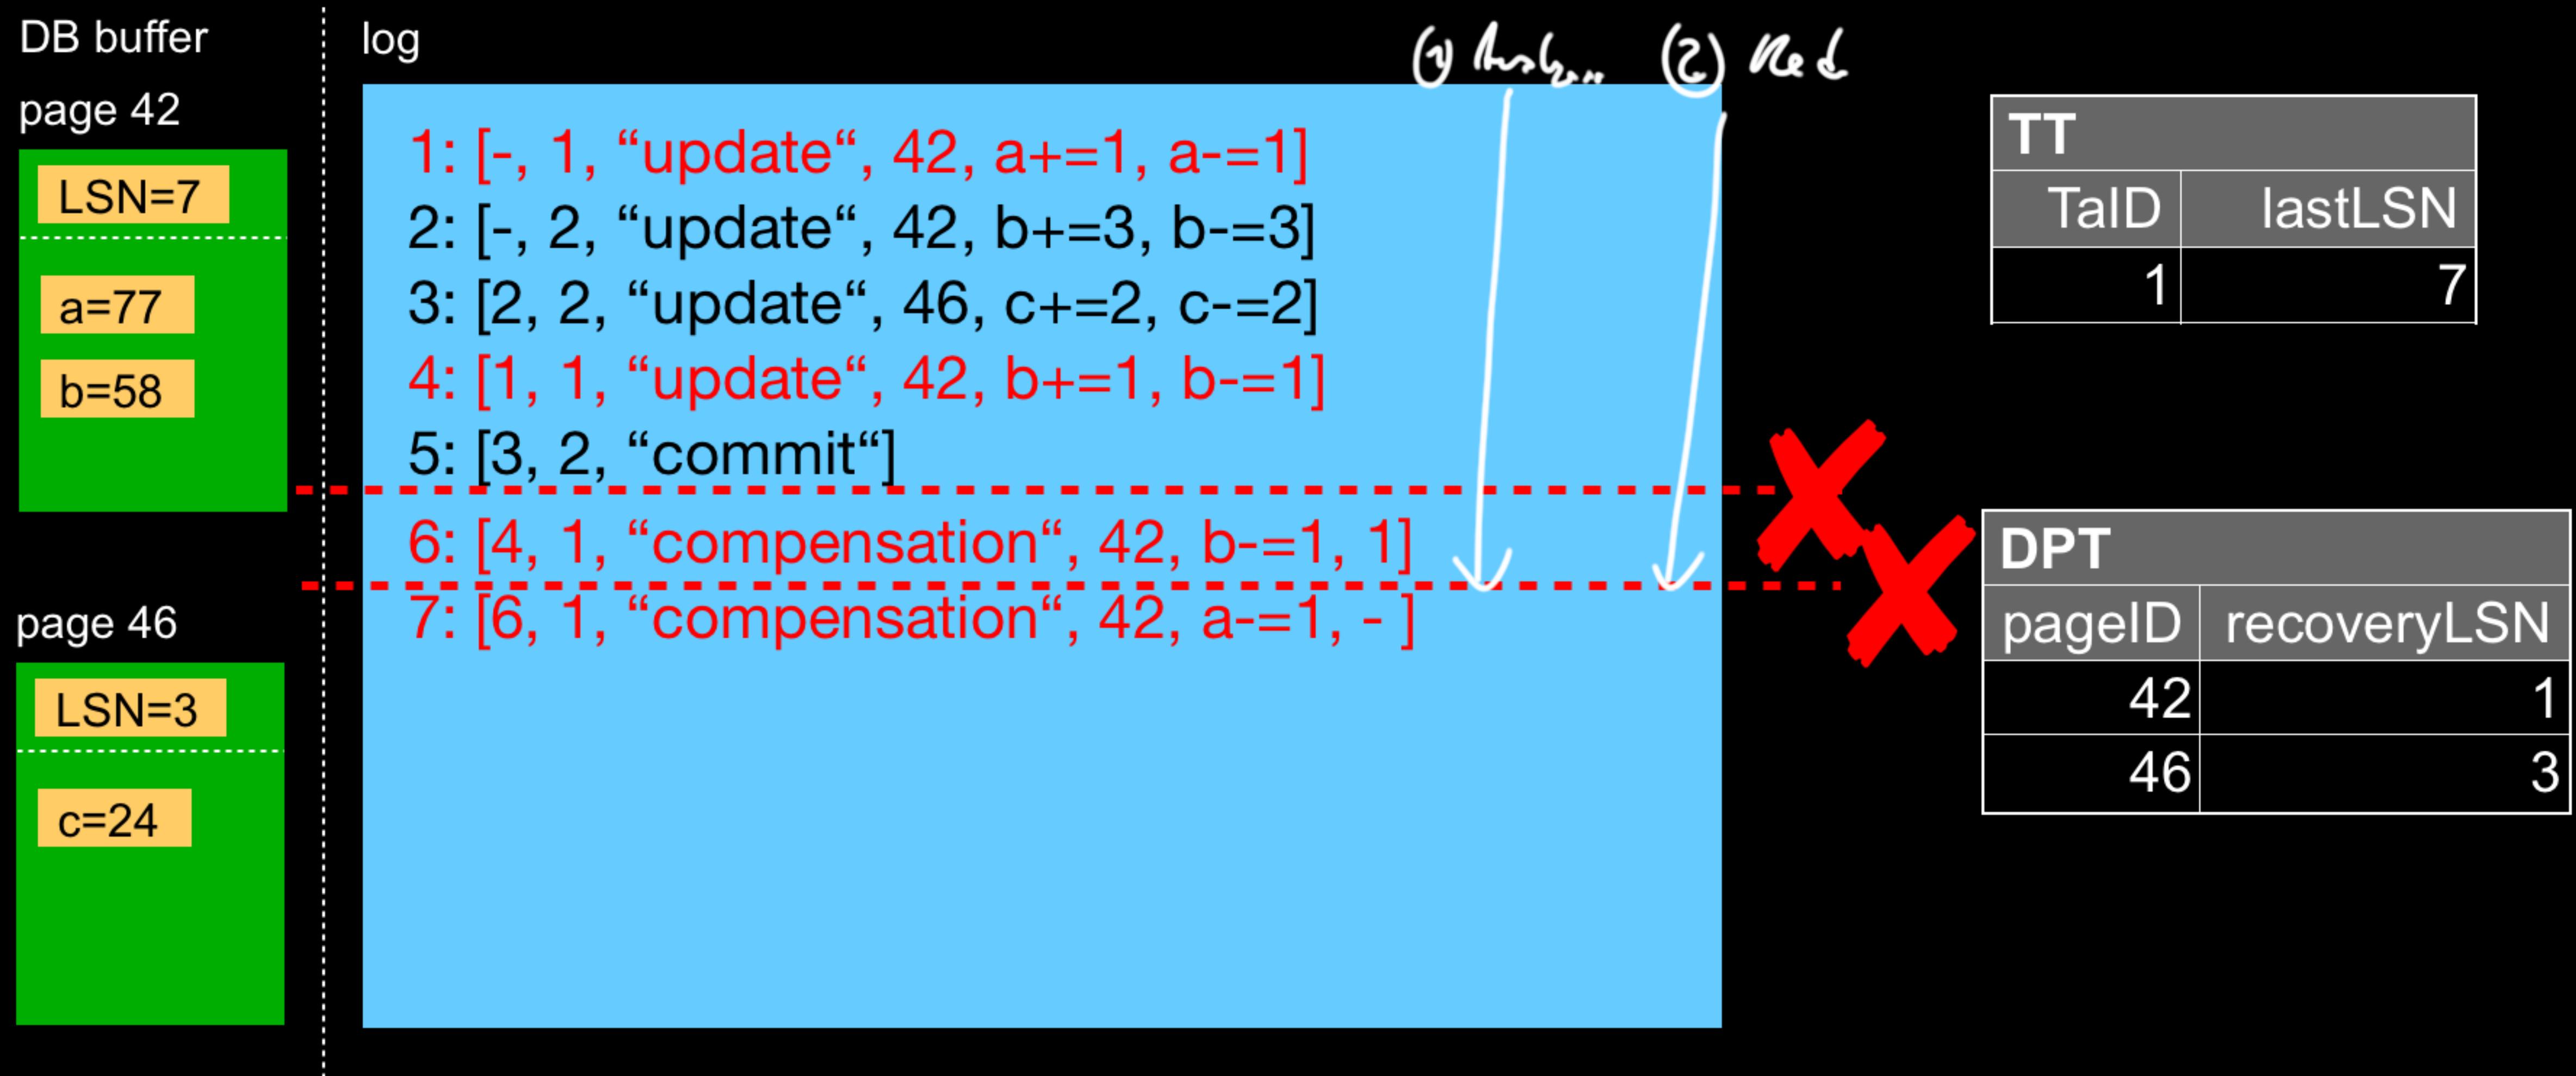
LSN=3
c=24

TT	
TaID	lastLSN
1	6

DPT	
pageID	recoveryLSN
42	1
46	3

# Example Log File

LSN: [prevLSN, TaID, type] all  
LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info] “update“  
LSN: [prevLSN, TaID, “compensation“, redoTheUndo info, undoNextLSN] “compensation“



# Example Log File

LSN: [prevLSN, TaID, type] all  
LSN: [prevLSN, TaID, “update“, pageID, redo info, undo info] “update“  
LSN: [prevLSN, TaID, “compensation“, redoTheUndo info, undoNextLSN] “compensation“

DB buffer

page 42

LSN=7
a=77
b=58

log

- 1: [-, 1, “update“, 42, a+=1, a-=1]
- 2: [-, 2, “update“, 42, b+=3, b-=3]
- 3: [2, 2, “update“, 46, c+=2, c-=2]
- 4: [1, 1, “update“, 42, b+=1, b-=1]
- 5: [3, 2, “commit“]
- 6: [4, 1, “compensation“, 42, b-=1, 1]
- 7: [6, 1, “compensation“, 42, a-=1, -]

page 46

LSN=3
c=24

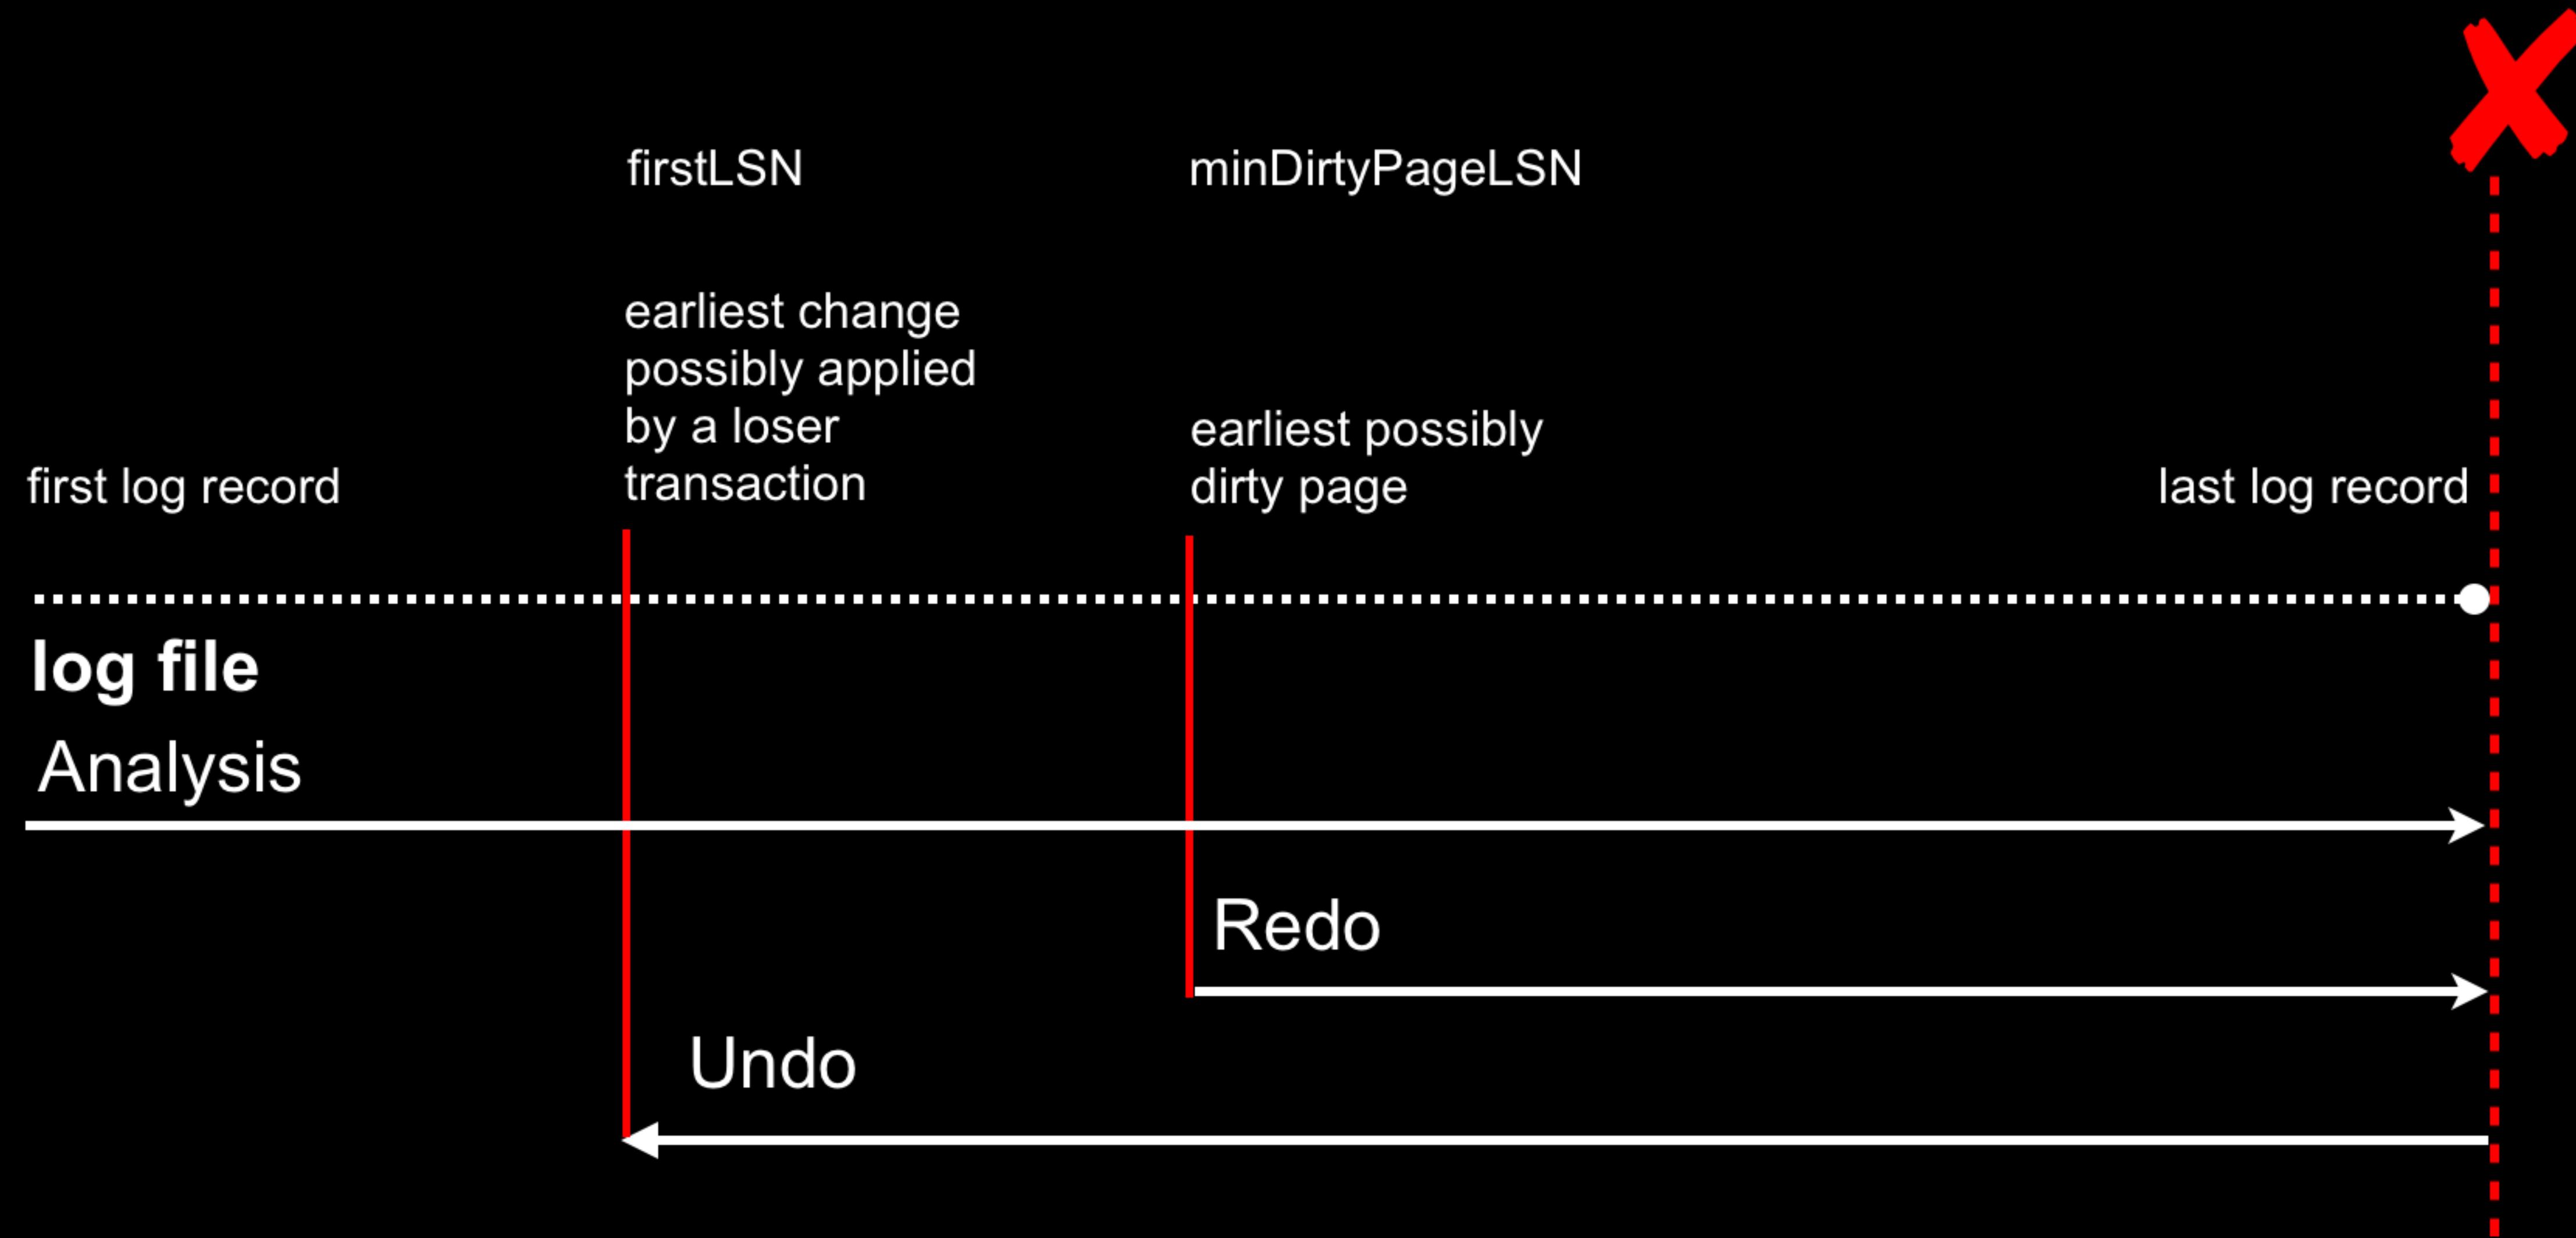
TT	
TaID	lastLSN

f:4KJN

DPT	
pageID	recoveryLSN
42	1
46	3

# Starting Points

```
minDirtyPageLSN := SELECT min(recoveryLSN) FROM DPT;
```

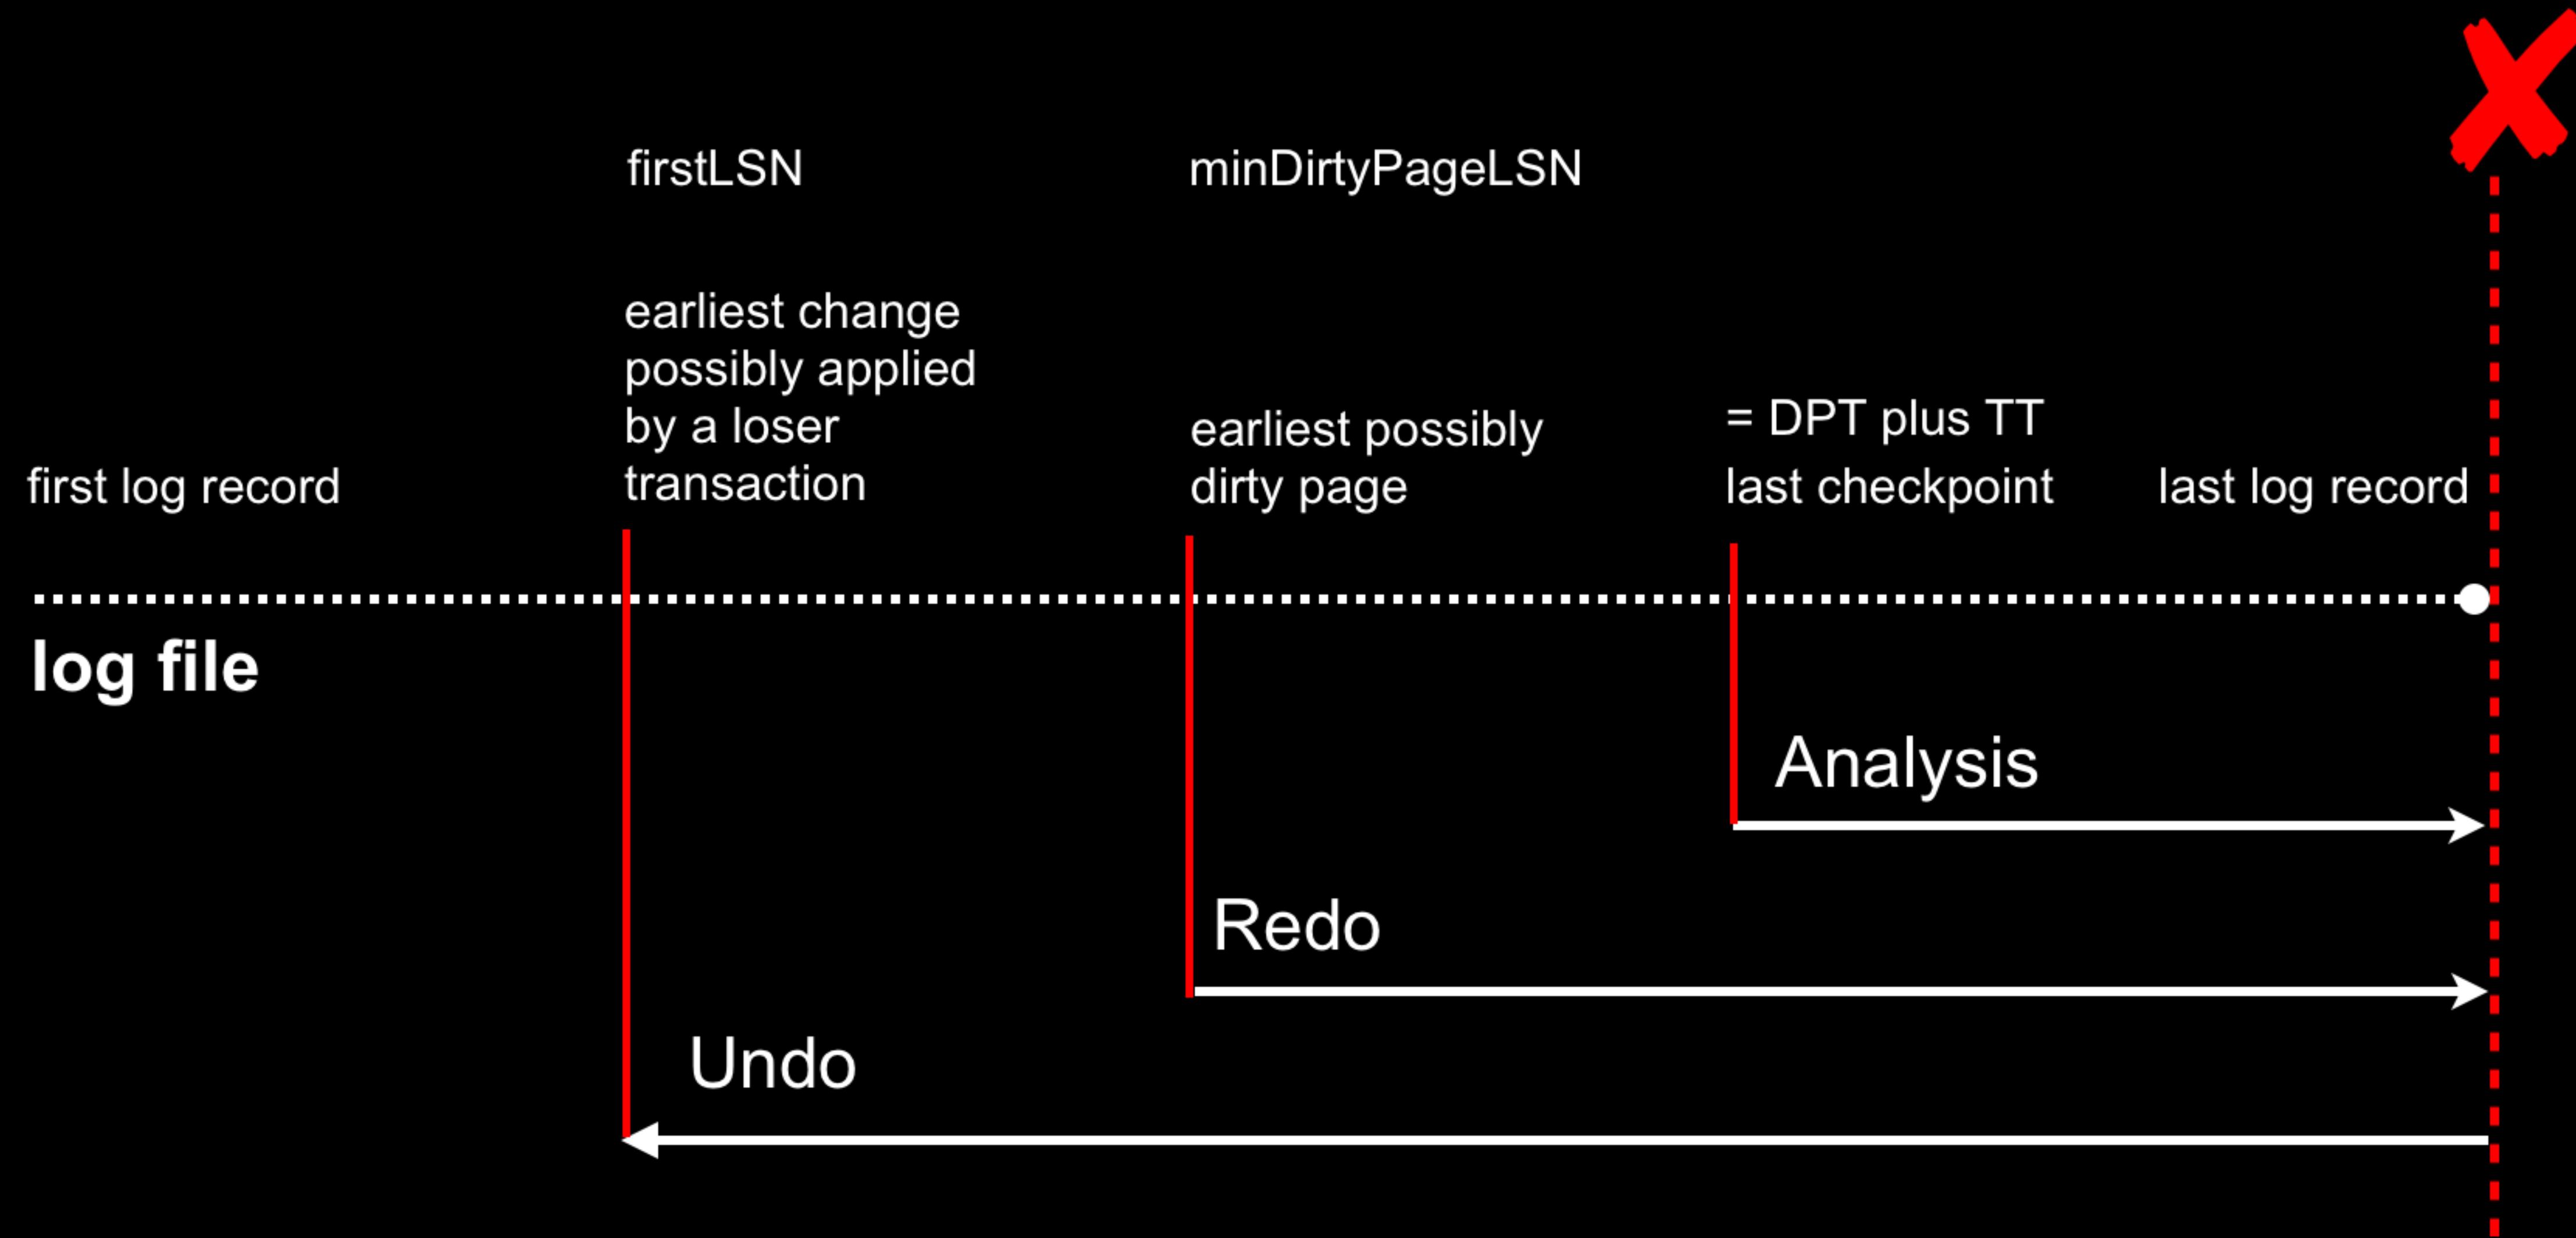


notice: for the moment no checkpointing!

note: relative starting points of the phases may differ

# Fuzzy Checkpoint

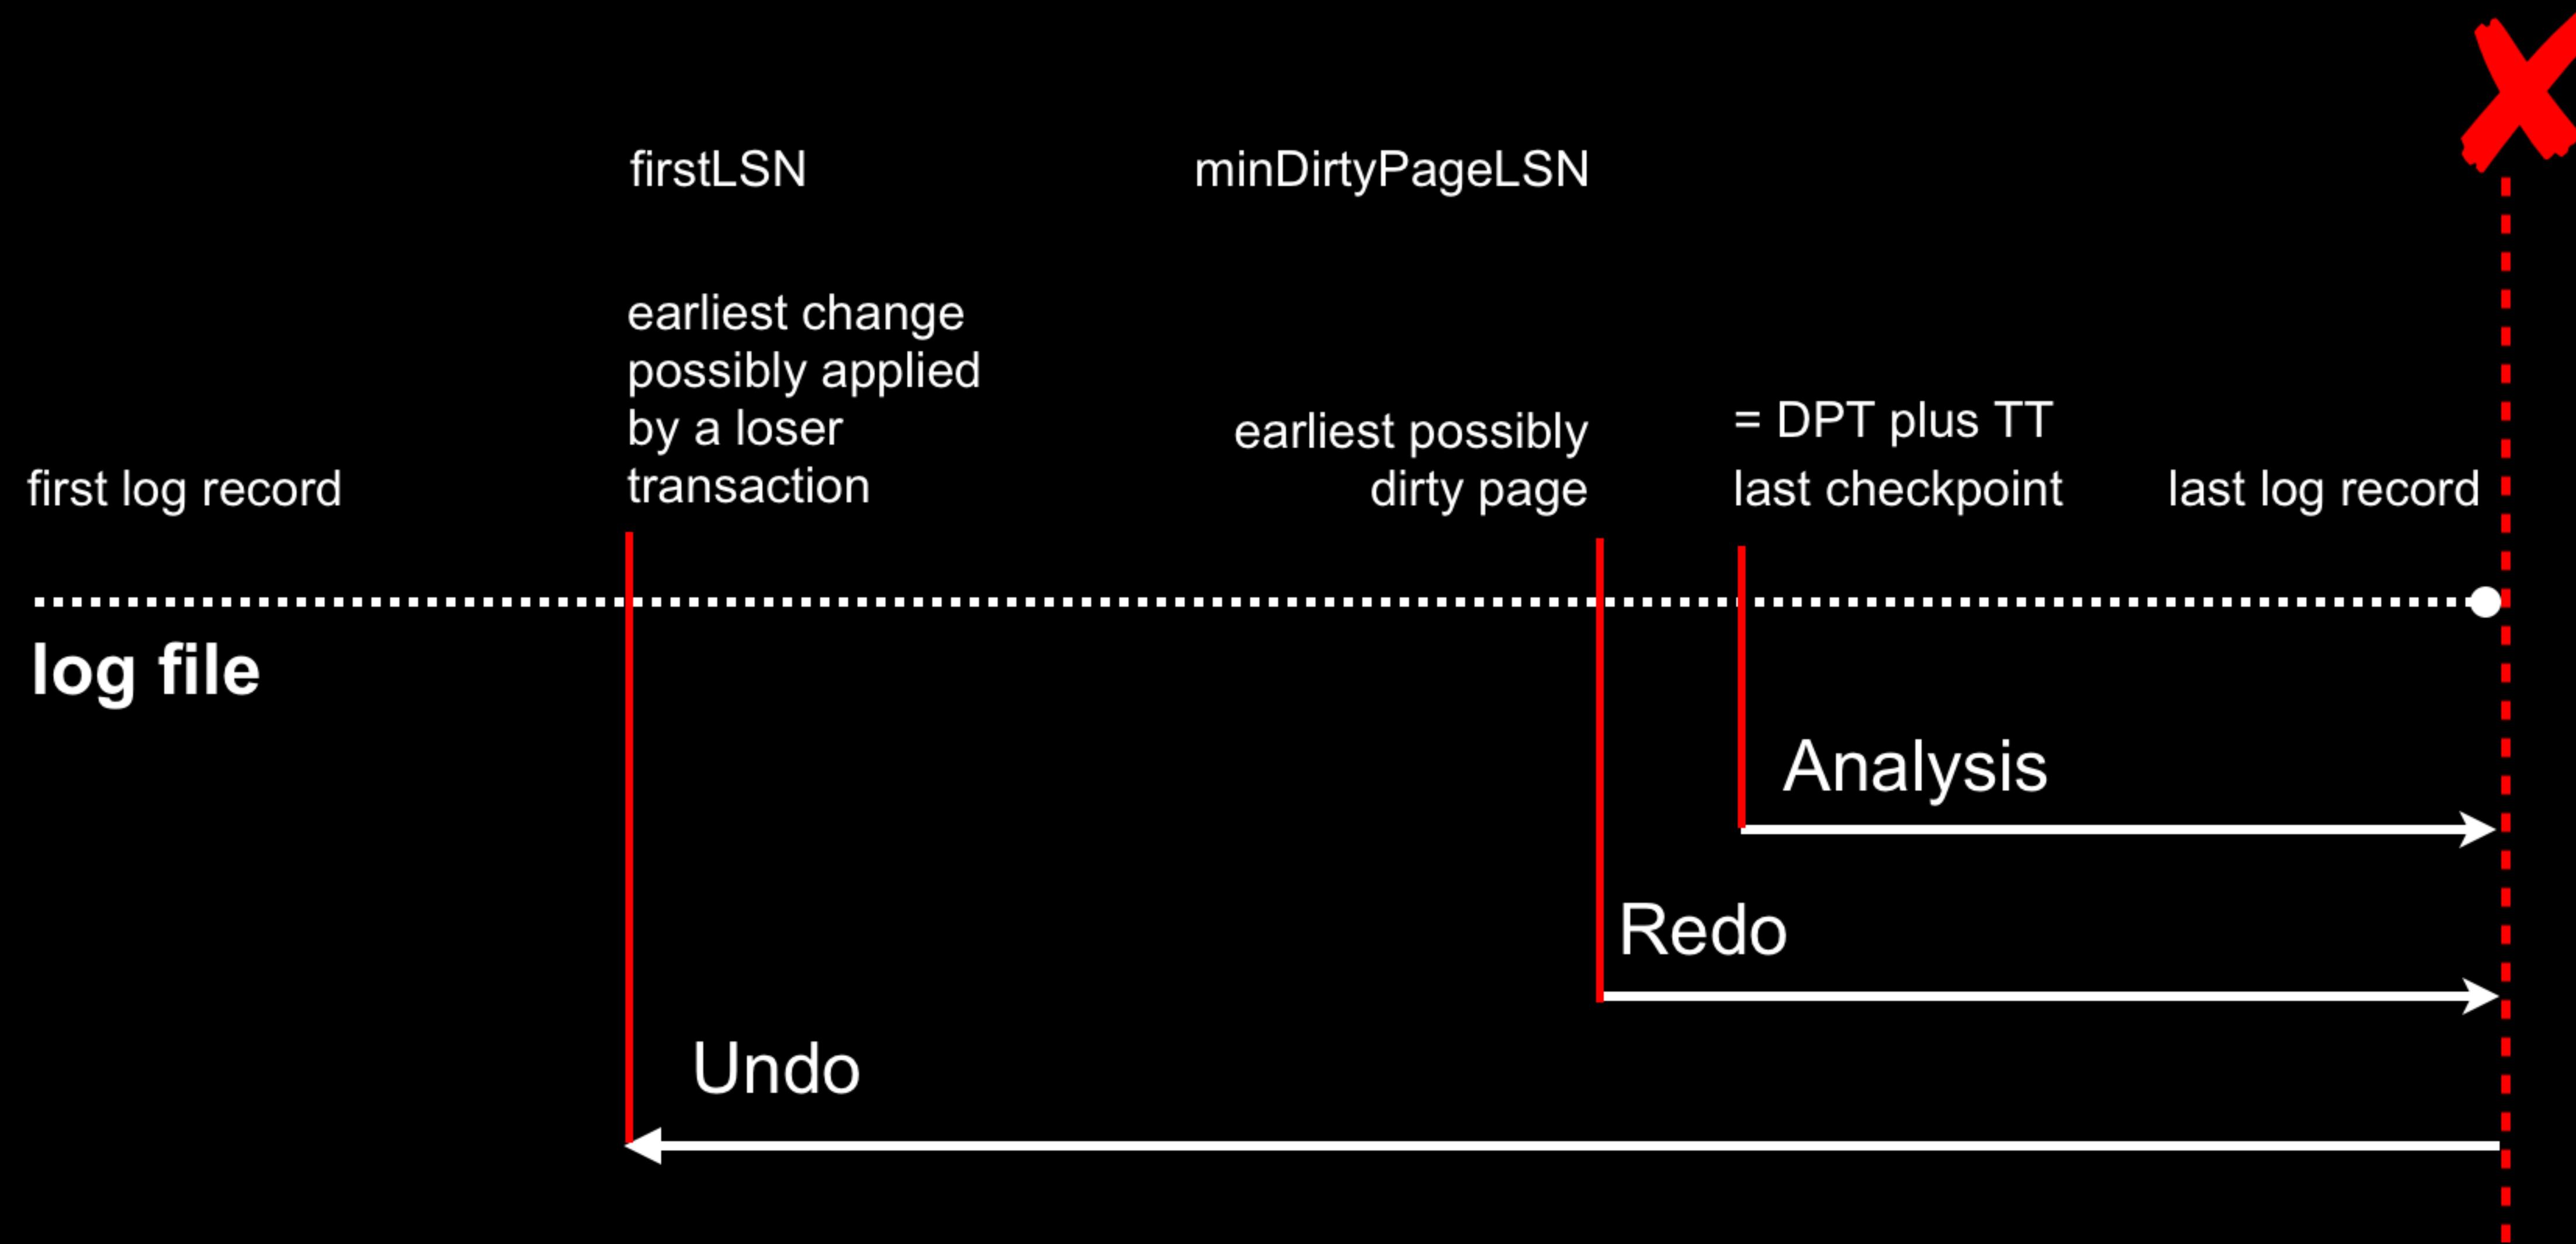
```
minDirtyPageLSN := SELECT min(recoveryLSN) FROM DPT;
```



note: relative starting points of the phases may differ

# Background Write Thread

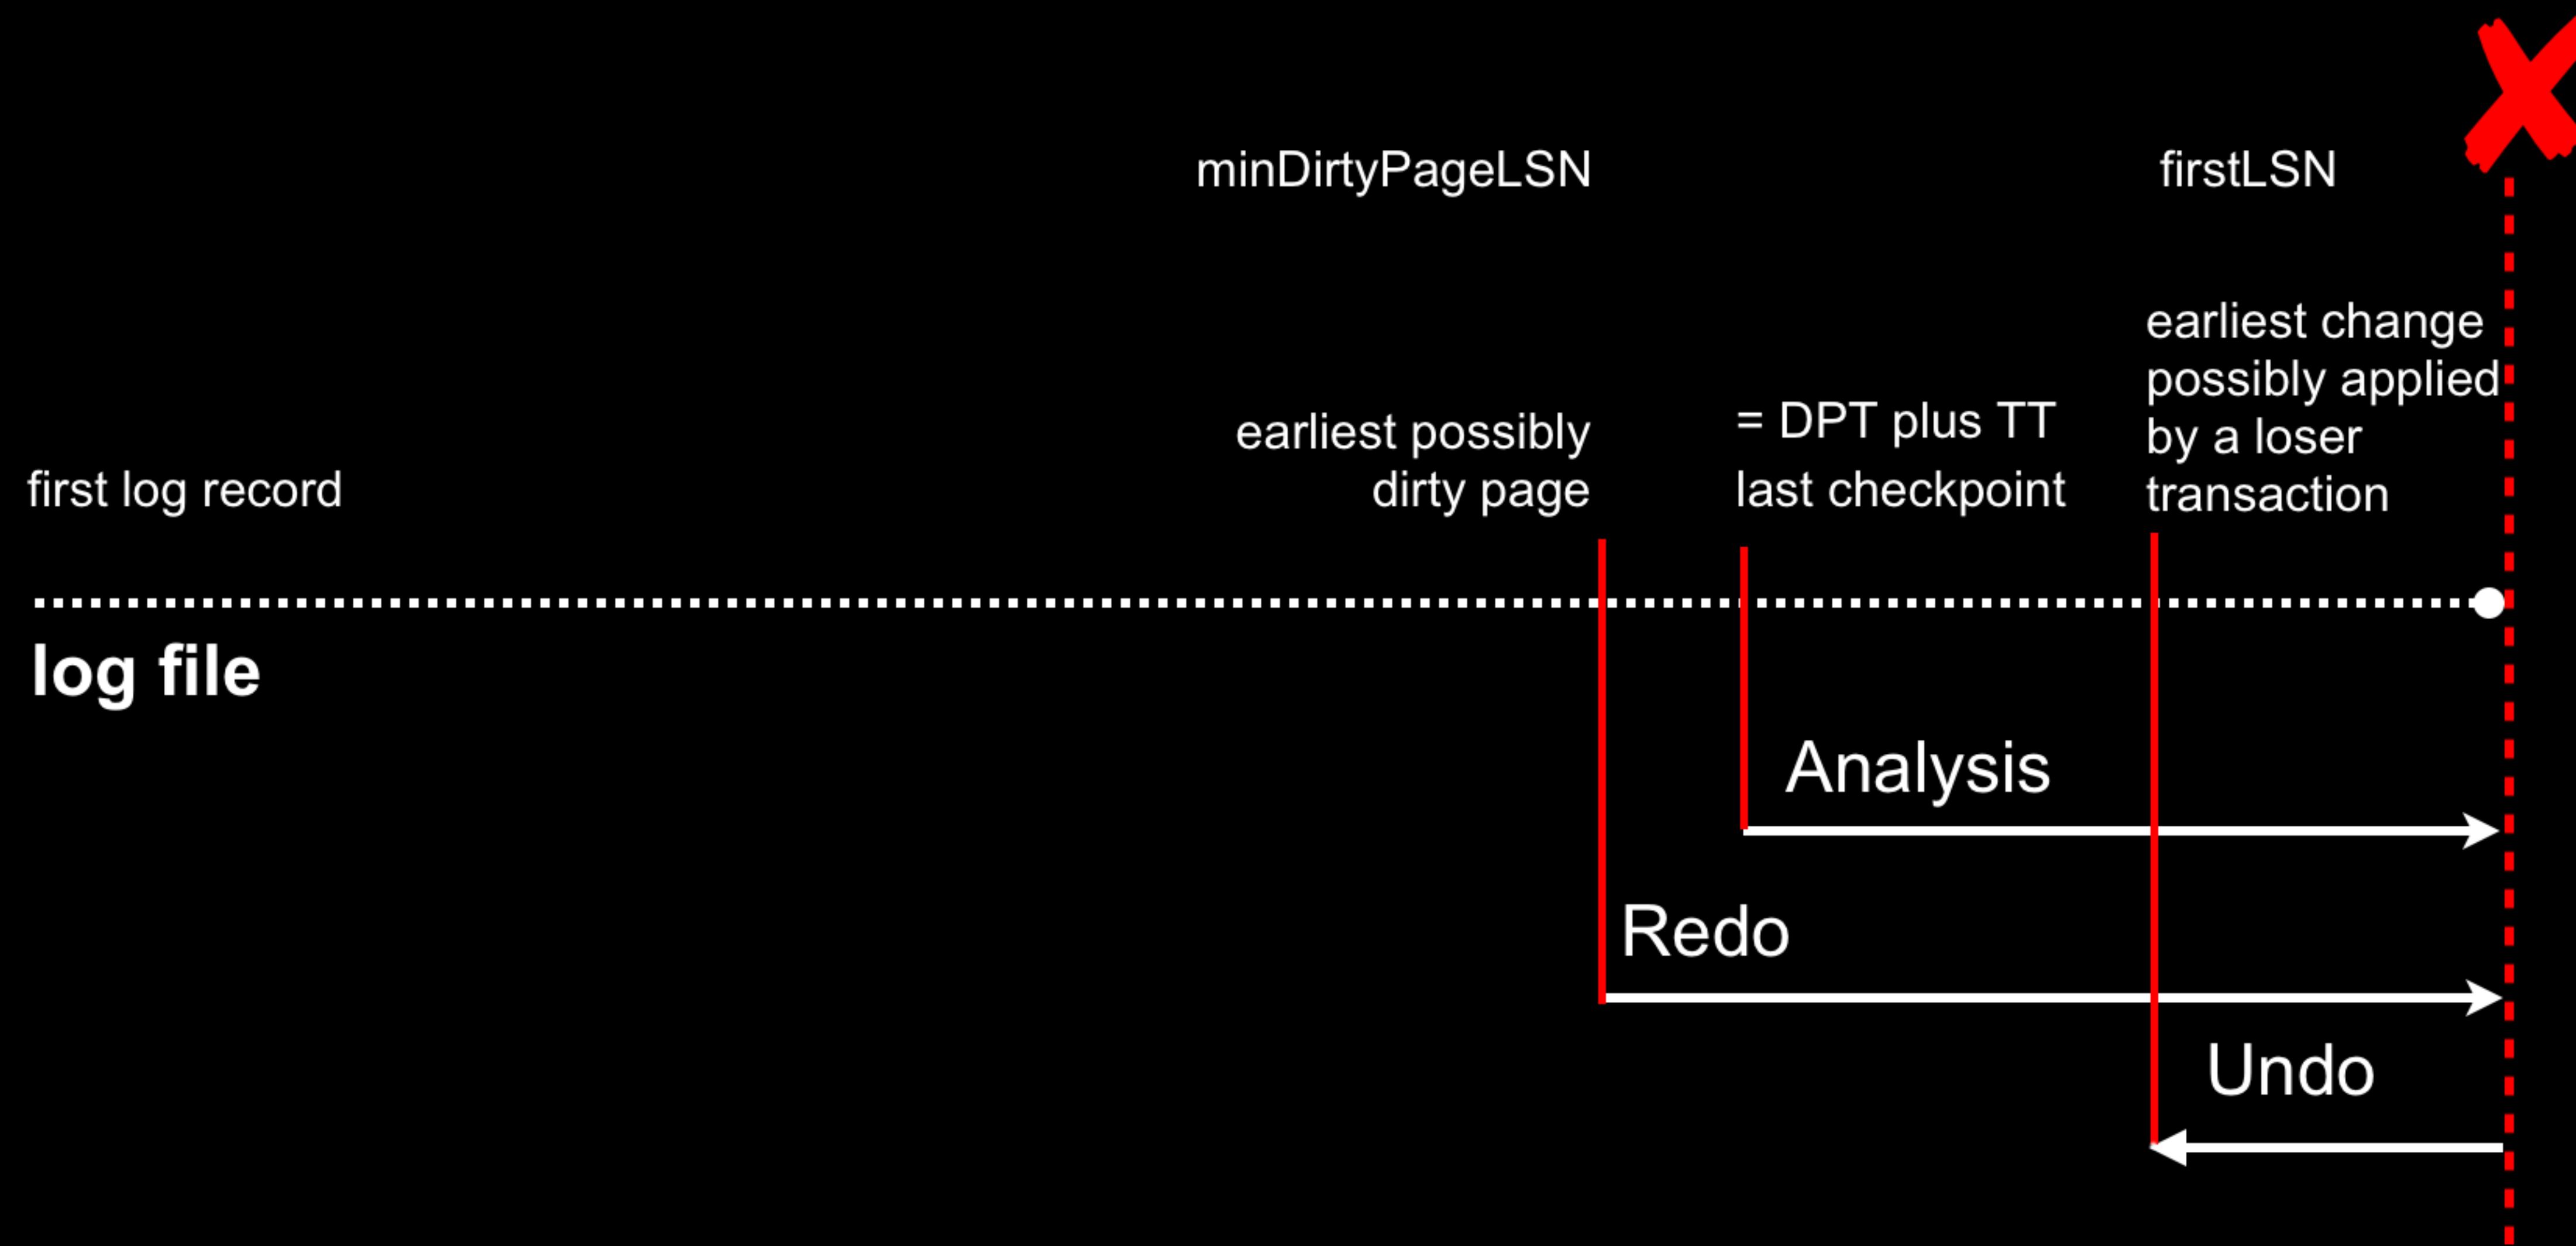
minDirtyPageLSN := `SELECT min(recoveryLSN) FROM DPT;`



note: relative starting points of the phases may differ

# Short Running Transactions

minDirtyPageLSN := `SELECT min(recoveryLSN) FROM DPT;`



note: relative starting points of the phases may differ

# Fuzzy Checkpointing Details

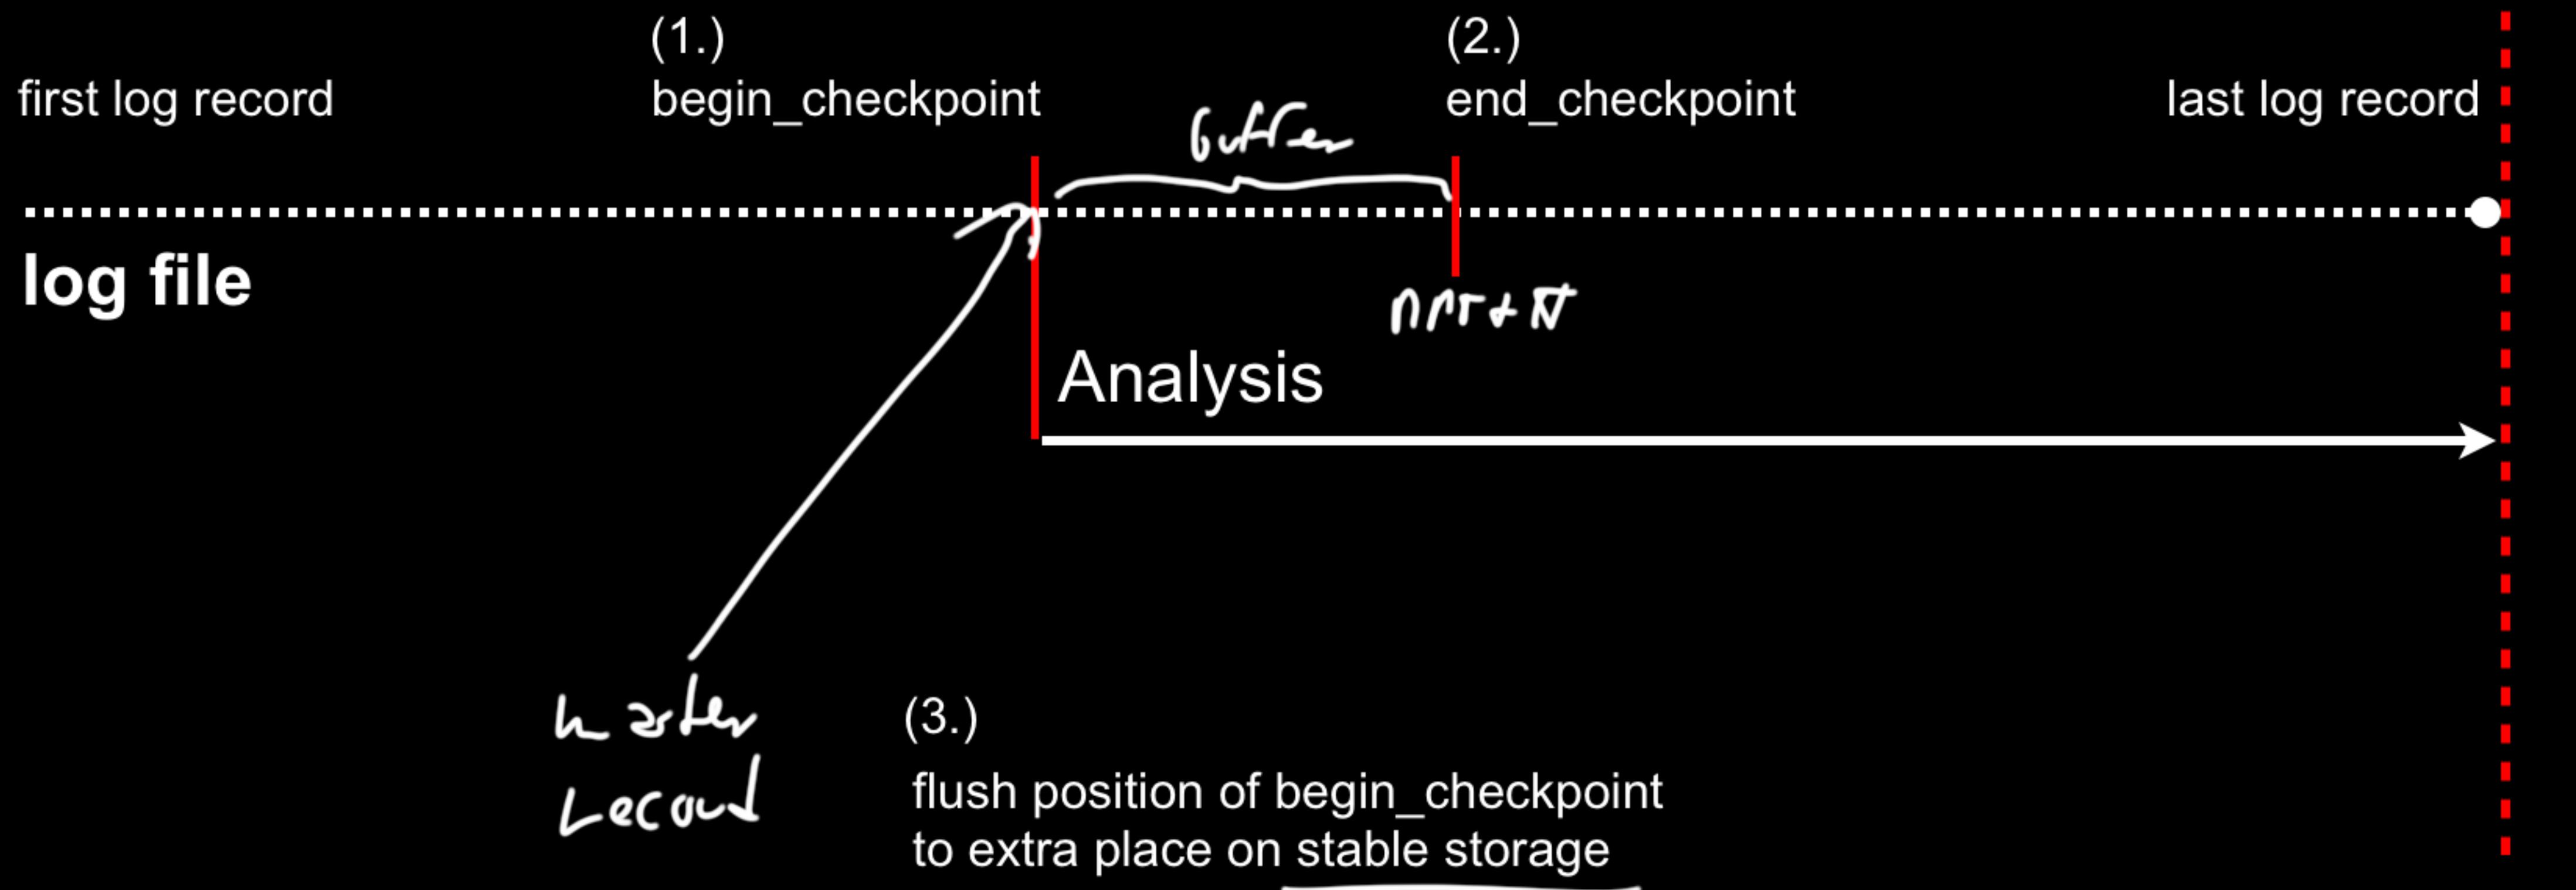
$\frac{V473}{TT + DPT}$ .  
25 of  
LSN 473

$\frac{V_{now}()}$   
TT + NPT  
25 of  
Low()

(COW)

LSN=473

= DPT plus TT  
as of begin\_checkpoint



# Tasks of the 3 Phases in ARIES

## Analysis Phase:

construct DPT and TT to reflect state at the time of the crash

compute minDirtyPageLSN (where to start Redo phase)

## Redo Phase:

Repeating history was done

repeat **all** actions: including those from loser transactions, including CLRs

# Tasks of the 3 Phases in ARIES

## Analysis Phase:

- construct DPT and TT to reflect state at the time of the crash

- compute minDirtyPageLSN (where to start Redo phase)

## Redo Phase:

- repeat **all** actions: including those from loser transactions, including CLRs

- restore database to where it was at the time of the crash

## Undo Phase:

- undo actions of all loser transactions

- log undos to CLRs