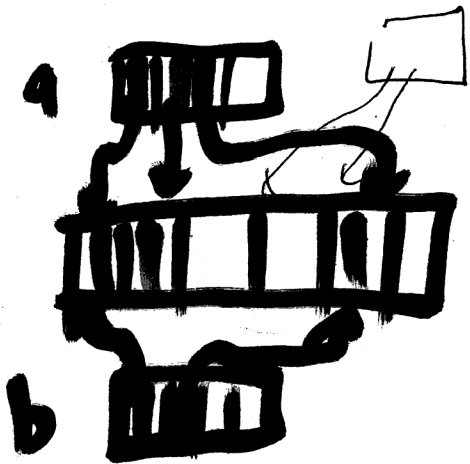




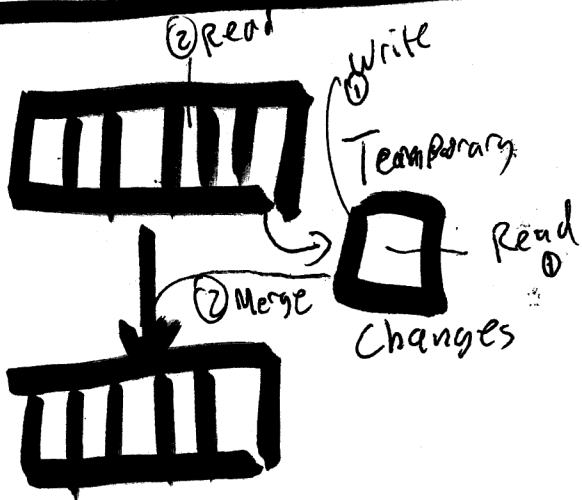
Twin Pages

Pro
Spatial Locality



Shadow Paging

Pro
Faster updates
Less Storage



Before T1
↓
Transaction Running
↓
Commit
↓
Merge
↓

→ Laser Marmot

OK

Can abort read by deleting changes

Naive Two-Phase Paging (Batched)

↳ Lose track of intermediate version

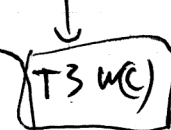
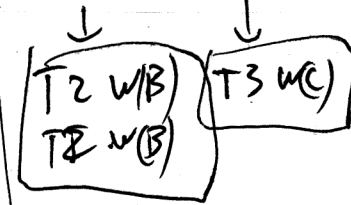
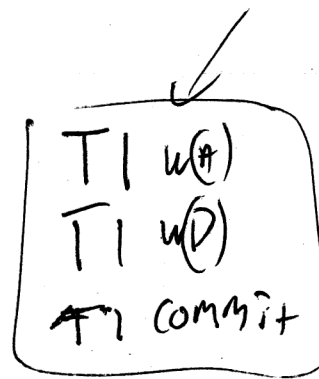
↳ Transactions delayed to end of batch

T1
↓ ← lost
T2
⋮

T1 writes page

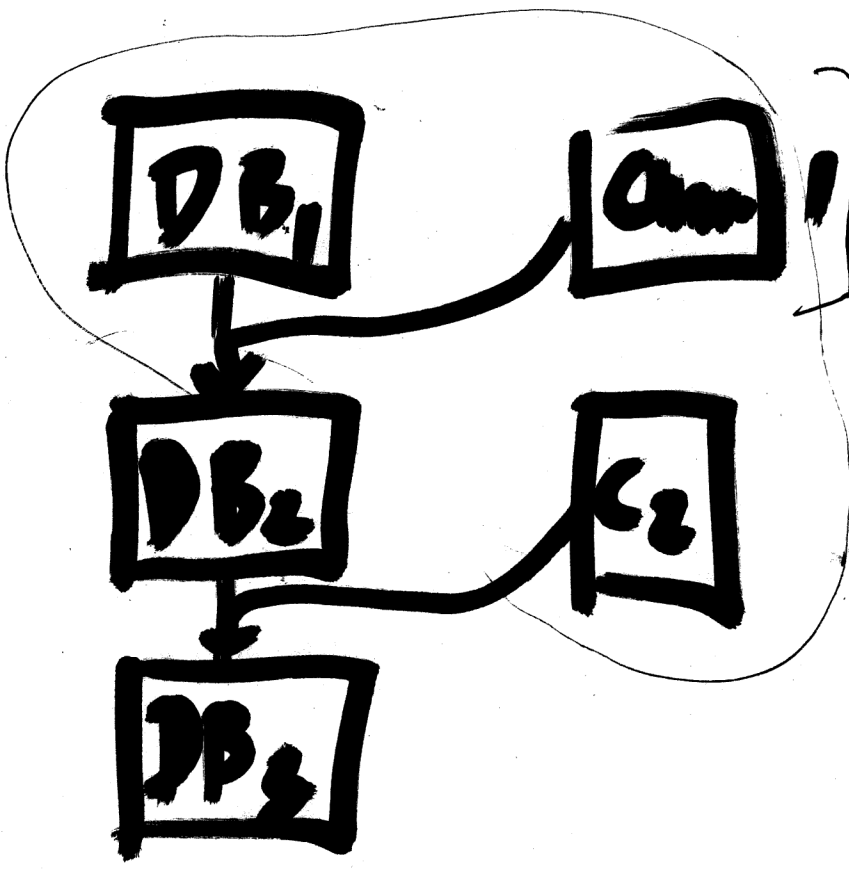
T1 W(A)
T2 W(B)
T3 W(C)
T2 W(B)
T1 W(D)
T1 COMMIT

Need to guarantee that this is safe on disk

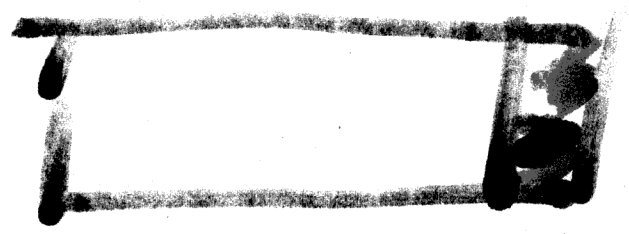


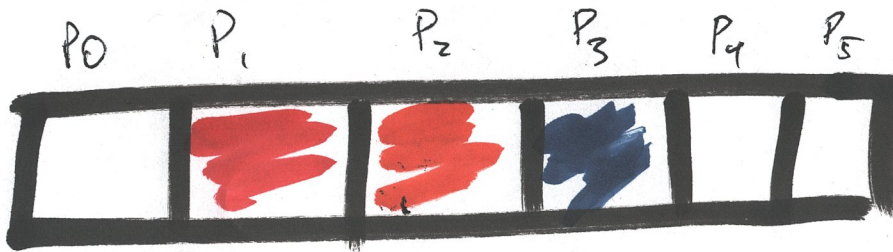
before we report T1 COMMITTED

fsync



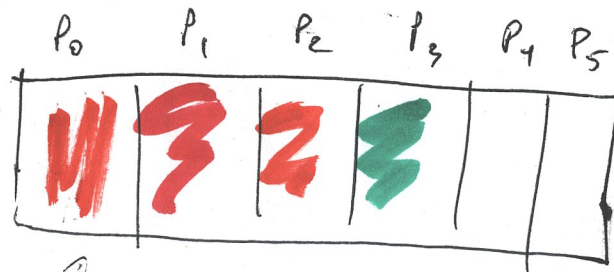
DB_1 exists $\leftarrow DB_1$
 \downarrow
 $DB_1 +$ Change₁ = DB_2 Changes exist $\leftarrow DB_1 +$ changes₁
 \downarrow
 DB_2 creation start = DB_2 $\leftarrow (DB_1 +$ changes₁)
 \downarrow
 DB_2 finishes + changes₂ $\leftarrow DB_2 +$ changes₂
 \downarrow
 $DB_1 +$ changes₁ + changes₂ = $DB_2 +$ changes₂





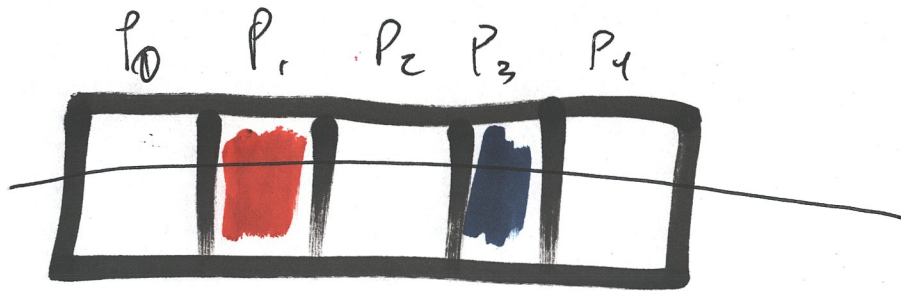
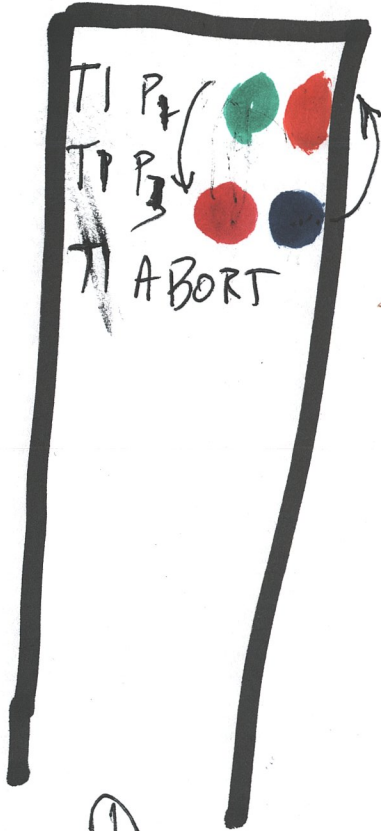
T1: P3
 T1: commit
 T2: P4
 T3: P0
 T2: Abort
 T3: commit

Differential Files

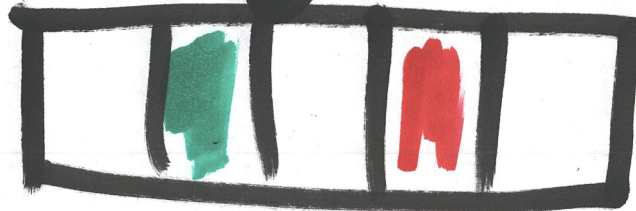


Merged
 as of T3

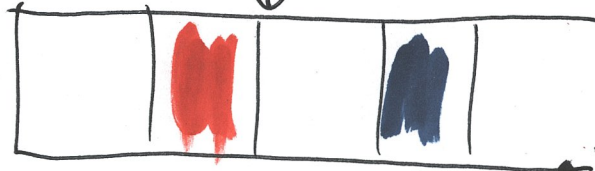
~~keep:~~



$T_1, w(P_1), w(P_3)$

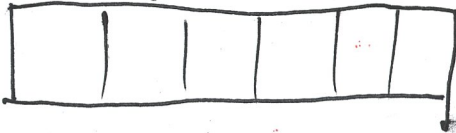


$T_1, ABORT$



time

Buffer Manager
(transient)



Disk
(persistent)



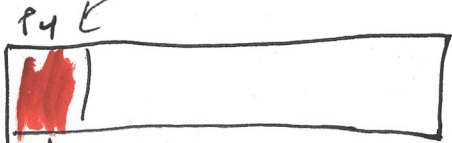
Log
(persistent)

Transaction Manager

T₁ = running

T₂ = running

Logs T₁ W(4)



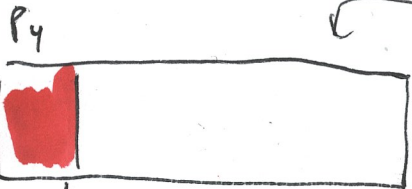
T₁ Commit



T₁ = Committed
T₂ = running



once on disk
we're good

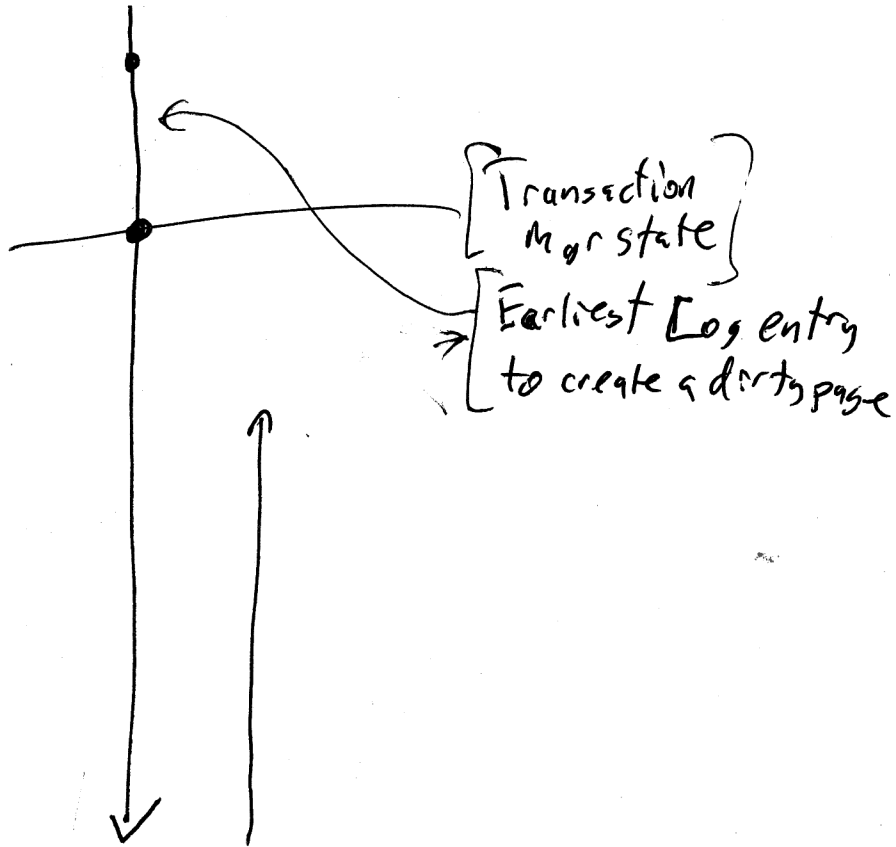


T₁ needs abort

T₁ = running
T₂ = commit
~~T₂ = ??~~

ok
because
T₂ never
wrote

Log



Replay Abort
running
acts