

Lecture 28:

Monday, December 9, 2002

1

Outline

- From the homework: Mr. Frumble's blues
- An exercise: counting the number of joins

- Redo logging – 17.3
- Redo/undo logging – 17.4

- Course evaluation forms

2

Understanding Hash Function Distribution

- $N = 100$ buckets
- Find the distribution of:
 $H('a00')$, $H('a01')$, ..., $H('a99')$
- $\text{Ascii}('a') = 97$, $\text{ascii}('0') = 48$
- Hence all values will start with:
 $(97+48+48) \bmod 100 = 93$
think of 93 as the new origin, and ignore it

3

Understanding Hash Function Distribution

- Hence the values of:
 $H('a00')$, $H('a01')$, ..., $H('a99')$
are:
 $0+0, 0+1, 0+2, \dots, 9+9$
- Observation 1: only buckets 0, 1, ..., 18 contain data !
- Observation 2:
 - Buckets 0 and 18 contain 1 data item
 - Buckets 1 and 17 contain 2 data items
 - ...
 - Bucket 9 contains 10 data items
- Then what happens with $H('a00000')$, ..., $H('a99999')$?

4

Counting the Number of Join Orders (Exercise)

$R_0(A_0, A_1) \bowtie R_1(A_1, A_2) \bowtie \dots \bowtie R_n(A_n, A_{n+1})$

- The number of left linear join trees is:
 $n!$
- The number of left linear join trees without cartesian products is:
 2^n (why ?)
- The number of bushy join trees is:
 $n!/(n+1) * C_n^{2n} = (2n)!/((n+1)*(n!))$
- The number of bushy join trees without cartesian product is:
 $2^{n-1}/(n+1) * C_n^{2n}$ (why ?)

5

Number of Subplans Inspected by Dynamic Programming

$R_0(A_0, A_1) \bowtie R_1(A_1, A_2) \bowtie \dots \bowtie R_n(A_n, A_{n+1})$

- The number of left linear subplans inspected is:
 $\sum_{k=1, n} C_n^k * k = n2^{n-1}$
- The number of left linear subplans without cartesian products inspected is:
 $\sum_{k=1, n} (n-k+1) * 2 = n(n+1)$ why ?
- The number of bushy join subplans inspected is:
 $\sum_{k=1, n} C_n^k * 2^k = 3^n$ why ?
- The number of bushy join subplans without cartesian product:
 $\sum_{k=1, n} (n-k+1) * (k-1) = n * n * (n-1) / 2 - n(n-1)(2n-1) / 6 = n(n-1)(n+1) / 6$

6

Redo Logging

Log records

- $\langle \text{START } T \rangle$ = transaction T has begun
- $\langle \text{COMMIT } T \rangle$ = T has committed
- $\langle \text{ABORT } T \rangle$ = T has aborted
- $\langle T, X, v \rangle$ = T has updated element X, and its new value is v

7

Redo-Logging Rules

R1: If T modifies X, then both $\langle T, X, v \rangle$ and $\langle \text{COMMIT } T \rangle$ must be written to disk before X is written to disk

- Hence: OUTPUTs are done late

8

Action	T	Mem A	Mem B	Disk A	Disk B	Log
						$\langle \text{START } T \rangle$
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	$\langle T, A, 16 \rangle$
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	$\langle T, B, 16 \rangle$
						$\langle \text{COMMIT } T \rangle$
OUTPUT(A)	16	16	16	16	8	
OUTPUT(B)	16	16	16	16	16	

9


Recovery with Redo Log

After system's crash, run recovery manager

- Step 1. Decide for each transaction T whether it is completed or not
 - $\langle \text{START } T \rangle \dots \langle \text{COMMIT } T \rangle \dots$ = yes
 - $\langle \text{START } T \rangle \dots \langle \text{ABORT } T \rangle \dots$ = yes
 - $\langle \text{START } T \rangle \dots$ = no
- Step 2. Read log from the beginning, redo all updates of committed transactions

10

Recovery with Redo Log



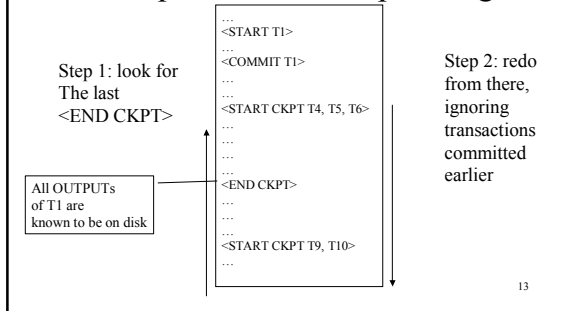
11

Nonquiescent Checkpointing

- Write a $\langle \text{START CKPT}(T1, \dots, Tk) \rangle$ where $T1, \dots, Tk$ are all active transactions
- Flush to disk all blocks of committed transactions (*dirty blocks*), while continuing normal operation
- When all blocks have been written, write $\langle \text{END CKPT} \rangle$

12

Redo Recovery with Nonquiescent Checkpointing



13

Comparison Undo/Redo

- Undo logging:
 - OUTPUT must be done early
 - If <COMMIT T> is seen, T definitely has written all its data to disk (hence, don't need to redo) – inefficient
- Redo logging
 - OUTPUT must be done late
 - If <COMMIT T> is not seen, T definitely has not written any of its data to disk (hence there is not dirty data on disk, no need to undo) – inflexible
- Would like more flexibility on when to OUTPUT: undo/redo logging (next)

14

Undo/Redo Logging

Log records, only one change

- <T,X,u,v>= T has updated element X, its old value was u, and its new value is v

15

Undo/Redo-Logging Rule

UR1: If T modifies X, then <T,X,u,v> must be written to disk before X is written to disk

Note: we are free to OUTPUT early or late (i.e. before or after <COMMIT T>)

16

Action	T	Mem A	Mem B	Disk A	Disk B	Log
						<START T>
READ(A,t)	8	8		8	8	
t:=t*2	16	8		8	8	
WRITE(A,t)	16	16		8	8	<T,A,8,16>
READ(B,t)	8	16	8	8	8	
t:=t*2	16	16	8	8	8	
WRITE(B,t)	16	16	16	8	8	<T,B,8,16>
OUTPUT(A)	16	16	16	16	8	
						<COMMIT T>
OUTPUT(B)	16	16	16	16	16	

17

Recovery with Undo/Redo Log

After system's crash, run recovery manager

- Redo all committed transaction, top-down
- Undo all uncommitted transactions, bottom-up

18

Recovery with Redo Log

↓

<START T1>
<T1,X1,v1>
<START T2>
<T2, X2, v2>
<START T3>
<T1,X3,v3>
<COMMIT T2>
<T3,X4,v4>
<T1,X5,v5>
...
...

↑