

Chapter 12: Indexing and Hashing

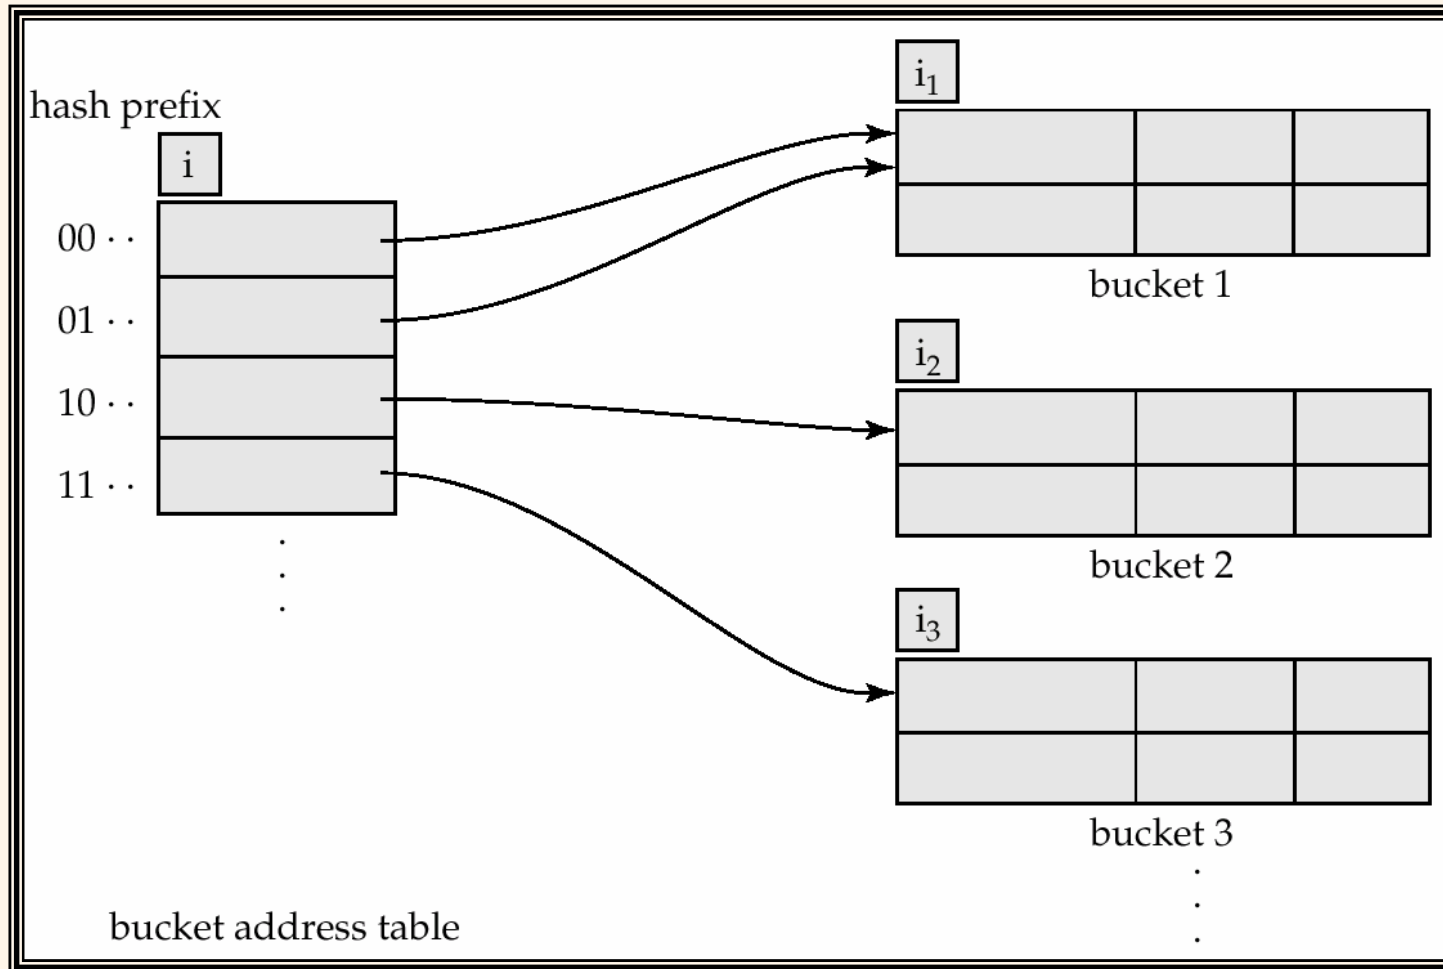
(Cnt.)

- ❖ Basic Concepts
- ❖ Ordered Indices
- ❖ B+-Tree Index Files
- ❖ B-Tree Index Files
- ❖ Static Hashing
- ❖ Dynamic Hashing
- ❖ Comparison of Ordered Indexing and Hashing
- ❖ Index Definition in SQL
- ❖ Multiple-Key Access

Dynamic Hashing

- ❖ Good for database that grows and shrinks in size
- ❖ Allows the hash function to be modified dynamically
- ❖ **Extendable hashing** – one form of dynamic hashing
 - Hash function generates values over a large range – typically b -bit integers, with $b = 32$.
 - At any time use only a prefix of the hash function to index into a table of bucket addresses.
 - Let the length of the prefix be i bits, $0 \leq i \leq 32$.
 - Bucket address table size = 2^i . Initially $i = 0$
 - Value of i grows and shrinks as the size of the database grows and shrinks.
 - Multiple entries in the bucket address table may point to a bucket.
 - Thus, actual number of buckets is $< 2^i$
 - The number of buckets also changes dynamically due to coalescing and splitting of buckets.

General Extendable Hash Structure




In this structure, $i_2 = i_3 = i$, whereas $i_1 = i - 1$ (see next slide for details)

Use of Extendable Hash Structure: Example




<i>branch-name</i>	<i>h(branch-name)</i>
Brighton	0010 1101 1111 1011 0010 1100 0011 0000
Downtown	1010 0011 1010 0000 1100 0110 1001 1111
Mianus	1100 0111 1110 1101 1011 1111 0011 1010
Perryridge	1111 0001 0010 0100 1001 0011 0110 1101
Redwood	0011 0101 1010 0110 1100 1001 1110 1011
Round Hill	1101 1000 0011 1111 1001 1100 0000 0001



Initial Hash structure, bucket size = 2

Example (Cont.)

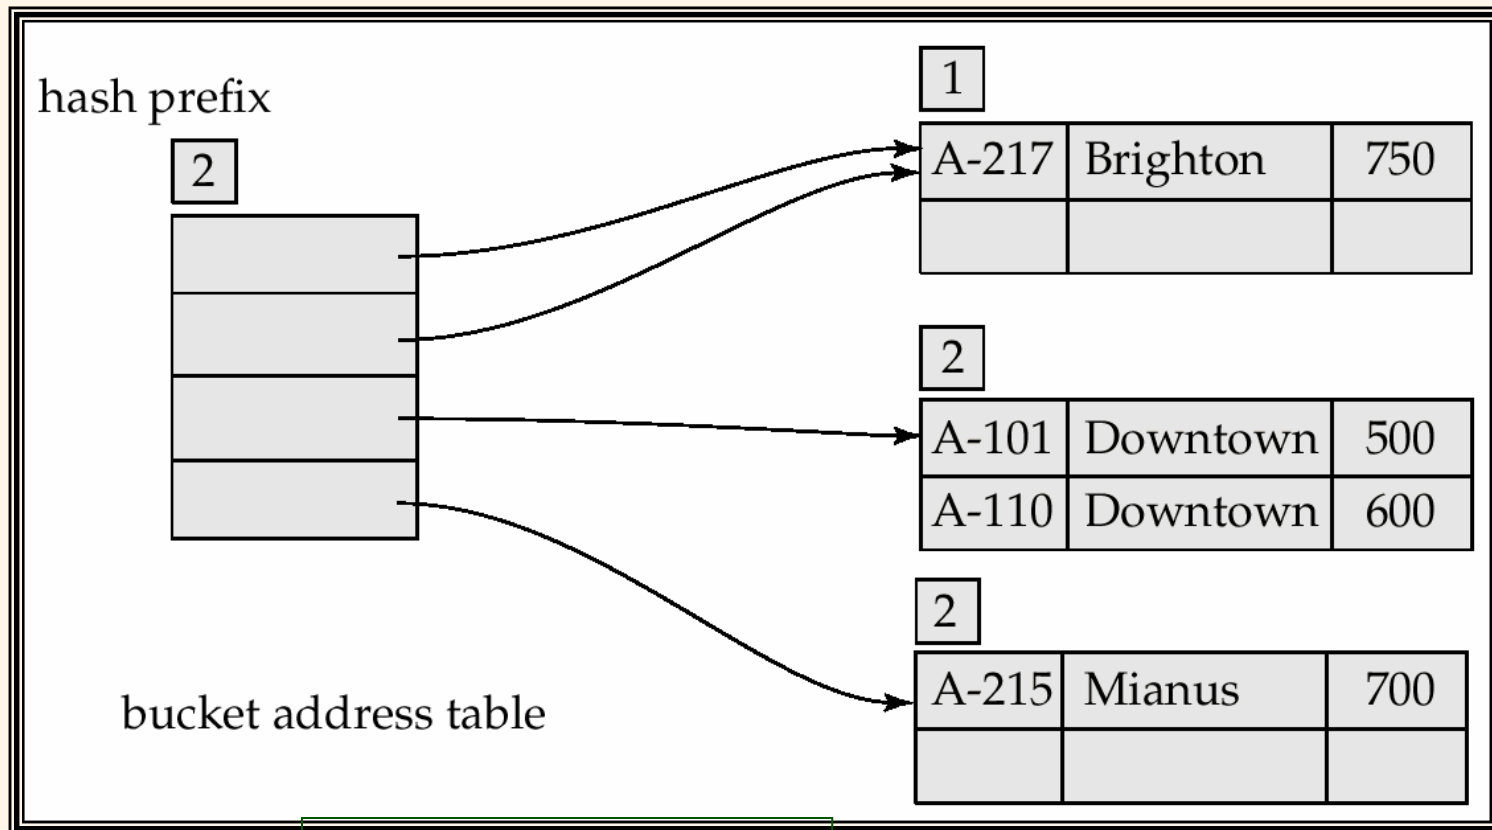
- ❖ Hash structure after insertion of one Brighton and two Downtown records



<i>Brighton</i>	<i>0010</i>
<i>Downtown</i>	<i>1010</i>

Example (Cont.)

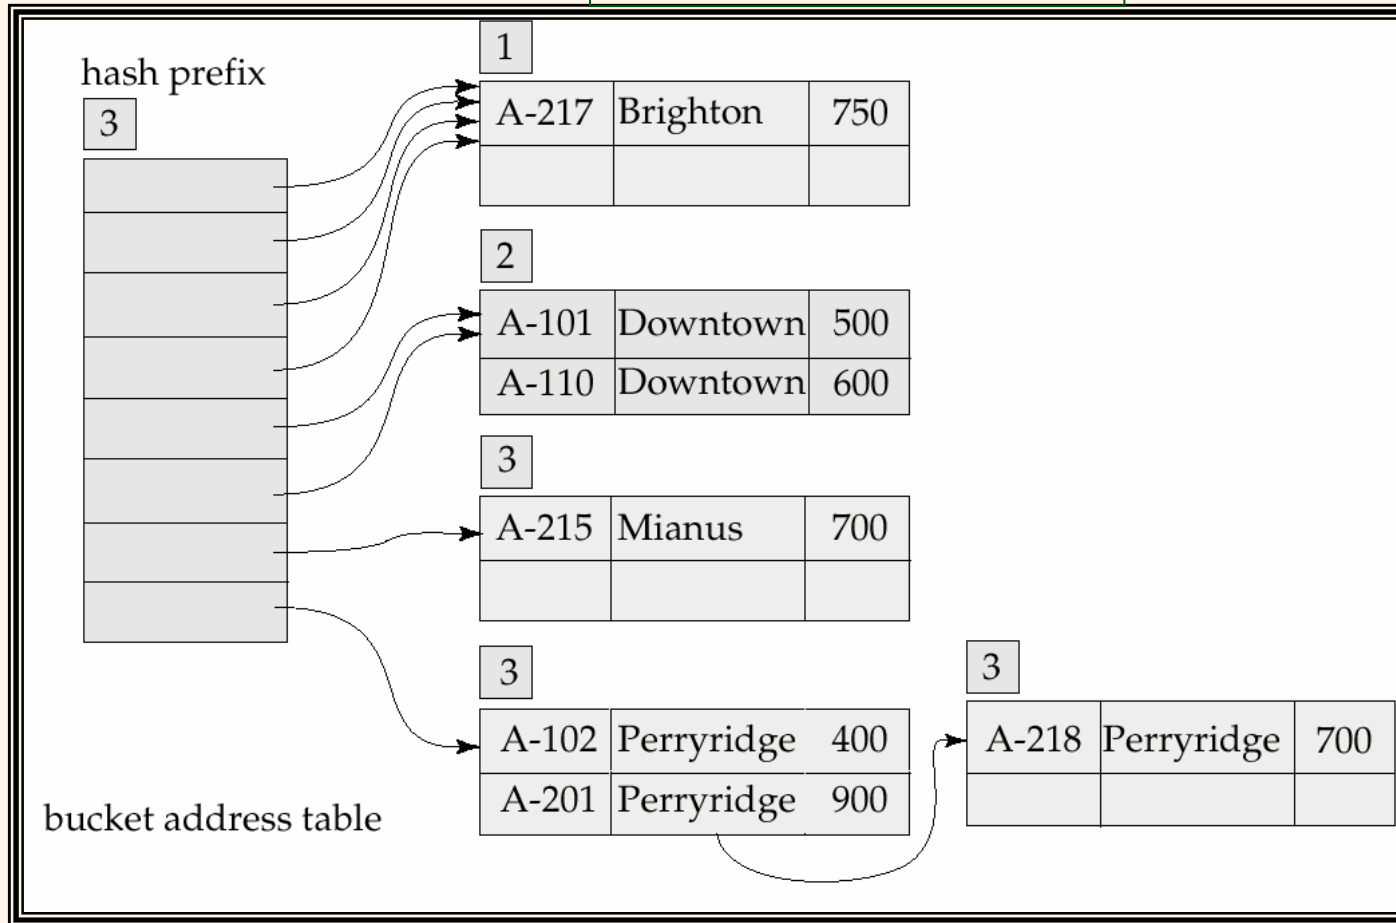

Hash structure after insertion of Mianus record



<i>Brighton</i>	<i>0010</i>
<i>Downtown</i>	<i>1010</i>
<i>Mianus</i>	<i>1100</i>

Example (Cont.)

Brighton	0010
Downtown	1010
Mianus	1100
Perryridge	1111



Hash structure after insertion of three Perryridge records

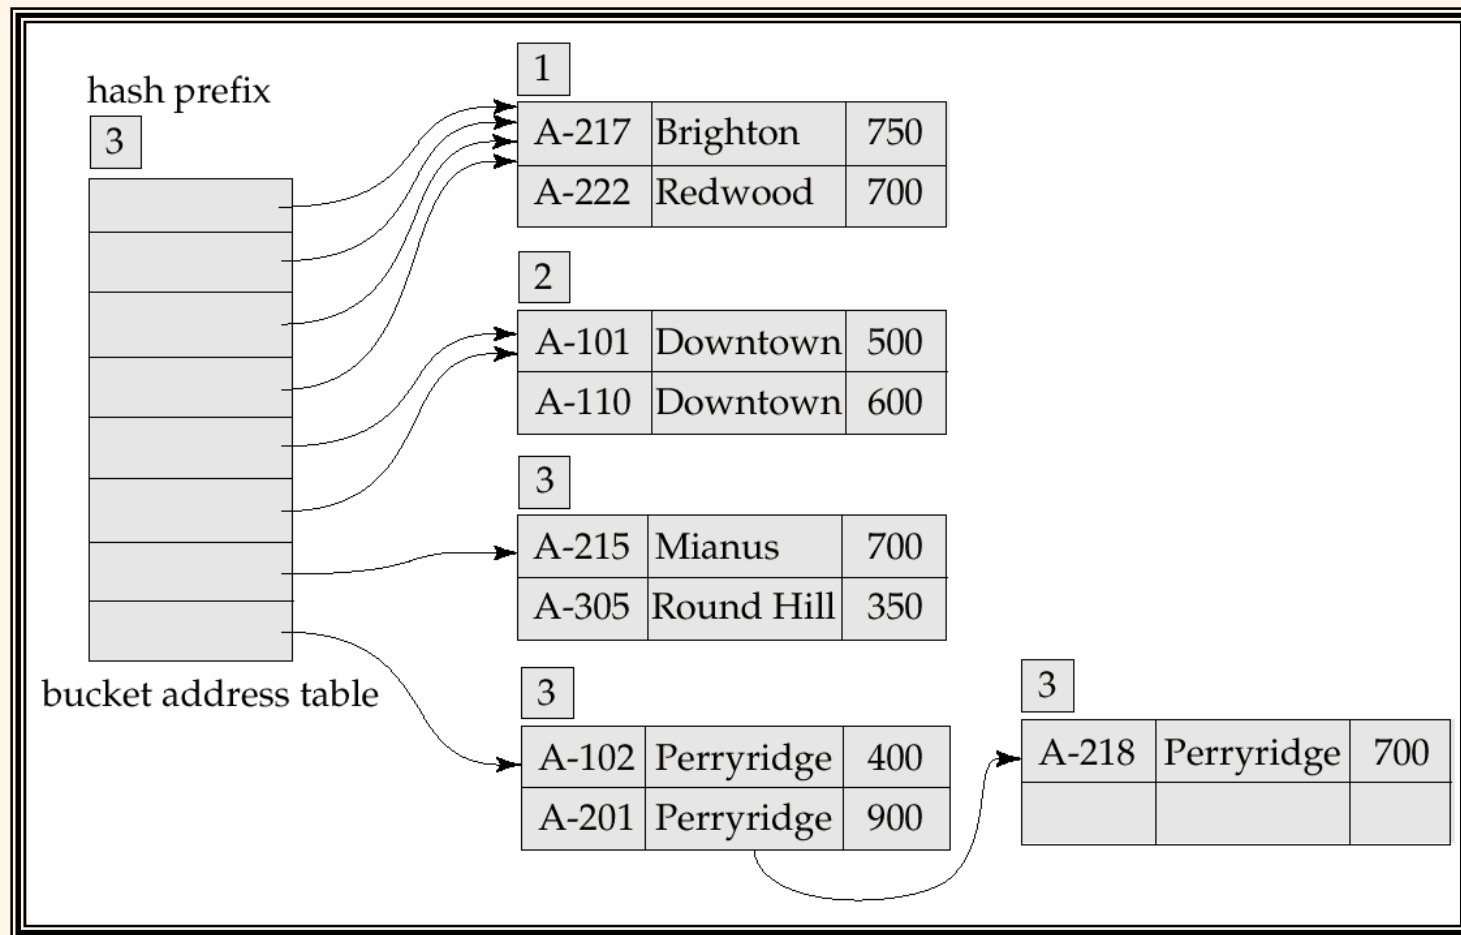
Example (Cont.)

Brighton	0010
Downtown	1010
Mianus	1100
Perryridge	1111

Redwood	0011
Round Hill	1101



- ❖ Hash structure after insertion of Redwood and Round Hill records



Use of Extendable Hash Structure

- ❖ Each bucket j stores a value i_j ; all the entries that point to the same bucket have the same values on the first i_j bits.
- ❖ To locate the bucket containing search-key K_j :
 1. Compute $h(K_j) = X$
 2. Use the first i high order bits of X as a displacement into bucket address table, and follow the pointer to appropriate bucket
- ❖ To insert a record with search-key value K_j
 - follow same procedure as look-up and locate the bucket, say j .
 - If there is room in the bucket j insert record in the bucket.
 - Else the bucket must be split and insertion re-attempted (next slide.)
 - Overflow buckets used instead in some cases (as the case for Perryridge in previous example)

Updates in Extendable Hash Structure



To split a bucket j when inserting record with search-key value K_j :

- ❖ If $i > i_j$ (more than one pointer to bucket j)
 - allocate a new bucket z , and set i_j and i_z to the old $i_j - + 1$.
 - make the second half of the bucket address table entries pointing to j to point to z
 - remove and reinsert each record in bucket j .
 - recompute new bucket for K_j and insert record in the bucket (further splitting is required if the bucket is still full)
- ❖ If $i = i_j$ (only one pointer to bucket j)
 - increment i and double the size of the bucket address table.
 - replace each entry in the table by two entries that point to the same bucket.
 - recompute new bucket address table entry for K_j
Now $i > i_j$ so use the first case above.

Updates in Extendable Hash Structure (Cont.)

- ❖ When inserting a value, if the bucket is full after several splits (that is, i reaches some limit b) create an overflow bucket instead of splitting bucket entry table further.
- ❖ To delete a key value,
 - locate it in its bucket and remove it.
 - The bucket itself can be removed if it becomes empty (with appropriate updates to the bucket address table).
 - Coalescing of buckets can be done (can coalesce only with a “buddy” bucket having same value of i_j and same $i_j - 1$ prefix, if it is present)
 - Decreasing bucket address table size is also possible
 - Note: decreasing bucket address table size is an expensive operation and should be done only if number of buckets becomes much smaller than the size of the table

Extendable Hashing vs. Other Schemes



- ❖ Benefits of extendable hashing:
 - Hash performance does not degrade with growth of file
 - Minimal space overhead
- ❖ Disadvantages of extendable hashing
 - Extra level of indirection to find desired record
 - Bucket address table may itself become very big (larger than memory)
 - Need a tree structure to locate desired record in the structure!
 - Changing size of bucket address table is an expensive operation
- ❖ Linear hashing is an alternative mechanism which avoids these disadvantages at the possible cost of more bucket overflows

Comparison of Ordered Indexing and Hashing



- ❖ Cost of periodic re-organization
- ❖ Relative frequency of insertions and deletions
- ❖ Is it desirable to optimize average access time at the expense of worst-case access time?
- ❖ Expected type of queries:
 - Hashing is generally better at retrieving records having a specified value of the key.
 - If range queries are common, ordered indices are to be preferred

Index Definition in SQL

- ❖ Create an index

create index <index-name> **on** <relation-name>
<attribute-list>)

E.g.: **create index** *b-index* **on** *branch*(*branch-name*)

- ❖ Use **create unique index** to indirectly specify and enforce the condition that the search key is a candidate key.

- ❖ To drop an index

drop index <index-name>

Multiple-Key Access

- ❖ Use multiple indices for certain types of queries.

- ❖ Example:

```
select account-number
```

```
from account
```

```
where branch-name = "Perryridge" and balance = 1000
```

- ❖ Possible strategies for processing query using indices on single attributes:

1. Use index on *branch-name* to find accounts with branch name of Perryridge; test *balance* = 1000.
2. Use index on *balance* to find accounts with balances of \$1000; test *branch-name* = "Perryridge".
3. Use *branch-name* index to find pointers to all records pertaining to the Perryridge branch. Similarly use index on *balance*. Take intersection of both sets of pointers obtained.

Indices on Multiple Attributes


Suppose we have an index on combined search-key
(*branch-name, balance*).

- ❖ With the **where** clause
where *branch-name* = “Perryridge” **and** *balance* = 1000
the index on the combined search-key will fetch only records that satisfy both conditions.
Using separate indices is less efficient – we may fetch many records (or pointers) that satisfy only one of the conditions.
- ❖ Can also efficiently handle
where *branch-name* = “Perryridge” **and** *balance* < 1000
- ❖ But cannot efficiently handle
where *branch-name* < “Perryridge” **and** *balance* = 1000
May fetch many records that satisfy the first but not the second condition.

Grid Files

- ❖ Structure used to speed the processing of general multiple search-key queries involving one or more comparison operators.
- ❖ The grid file has a single grid array and one linear scale for each search-key attribute. The grid array has number of dimensions equal to number of search-key attributes.
- ❖ Multiple cells of grid array can point to same bucket
- ❖ To find the bucket for a search-key value, locate the row and column of its cell using the linear scales and follow pointer

Example Grid File for account



Queries on a Grid File

- ❖ A grid file on two attributes A and B can handle queries of all following forms with reasonable efficiency
 - $(a_1 \leq A \leq a_2)$
 - $(b_1 \leq B \leq b_2)$
 - $(a_1 \leq A \leq a_2 \wedge b_1 \leq B \leq b_2),,$
- ❖ E.g., to answer $(a_1 \leq A \leq a_2 \wedge b_1 \leq B \leq b_2),$ use linear scales to find corresponding candidate grid array cells, and look up all the buckets pointed to from those cells.

Grid Files (Cont.)

- ❖ During insertion, if a bucket becomes full, new bucket can be created if more than one cell points to it.
 - Idea similar to extendable hashing, but on multiple dimensions
 - If only one cell points to it, either an overflow bucket must be created or the grid size must be increased
- ❖ Linear scales must be chosen to uniformly distribute records across cells.
 - Otherwise there will be too many overflow buckets.
- ❖ Periodic re-organization to increase grid size will help.
 - But reorganization can be very expensive.
- ❖ Space overhead of grid array can be high.
- ❖ R-trees (Chapter 23) are an alternative

Bitmap Indices

- ❖ Bitmap indices are a special type of index designed for efficient querying on multiple keys
- ❖ Records in a relation are assumed to be numbered sequentially from, say, 0
 - Given a number n it must be easy to retrieve record n
 - Particularly easy if records are of fixed size
- ❖ Applicable on attributes that take on a relatively small number of distinct values
 - E.g. gender, country, state, ...
 - E.g. income-level (income broken up into a small number of levels such as 0-9999, 10000-19999, 20000-50000, 50000- infinity)
- ❖ A bitmap is simply an array of bits

Bitmap Indices (Cont.)

❖ In its simplest form a bitmap index on an attribute has a bitmap for each value of the attribute

- Bitmap has as many bits as records
- In a bitmap for value v , the bit for a record is 1 if the record has the value v for the attribute, and is 0 otherwise

record number	<i>name</i>	<i>gender</i>	<i>address</i>	<i>income-level</i>	Bitmaps for <i>gender</i>		Bitmaps for <i>income-level</i>	
0	John	m	Perryridge	L1	m	1 0 0 1 0	L1	1 0 1 0 0
1	Diana	f	Brooklyn	L2	f	0 1 1 0 1	L2	0 1 0 0 0
2	Mary	f	Jonestown	L1			L3	0 0 0 0 1
3	Peter	m	Brooklyn	L4			L4	0 0 0 1 0
4	Kathy	f	Perryridge	L3			L5	0 0 0 0 0

Bitmap Indices (Cont.)

- ❖ Bitmap indices are useful for queries on multiple attributes
 - not particularly useful for single attribute queries
- ❖ Queries are answered using bitmap operations
 - Intersection (and)
 - Union (or)
 - Complementation (not)
- ❖ Each operation takes two bitmaps of the same size and applies the operation on corresponding bits to get the result bitmap
 - E.g. $100110 \text{ AND } 110011 = 100010$
 $100110 \text{ OR } 110011 = 110111$
 $\text{NOT } 100110 = 011001$
 - Males with income level L1: $10010 \text{ AND } 10100 = 10000$
 - Can then retrieve required tuples.
 - Counting number of matching tuples is even faster

Bitmap Indices (Cont.)

- ❖ Bitmap indices generally very small compared with relation size
 - E.g. if record is 100 bytes, space for a single bitmap is 1/800 of space used by relation.
 - If number of distinct attribute values is 8, bitmap is only 1% of relation size
- ❖ Deletion needs to be handled properly
 - Existence bitmap to note if there is a valid record at a record location
 - Needed for complementation
 - $\text{not}(A=v)$: $(\text{NOT } \text{bitmap-}A\text{-}v) \text{ AND } \text{ExistenceBitmap}$
- ❖ Should keep bitmaps for all values, even null value
 - To correctly handle SQL null semantics for $\text{NOT}(A=v)$:
 - intersect above result with $(\text{NOT } \text{bitmap-}A\text{-Null})$

End of Chapter

Sample account File



A-217	Brighton	750
A-101	Downtown	500
A-1 10	Downtown	600
A-215	Mianus	700
A-102	Perryridge	400
A-201	Perryridge	900
A-218	Perryridge	700
A-222	Redwood	700
A-305	Round Hill	350