

System Catalog

- A DBMS stores meta-data about databases in its internal catalog.
- List of tables, columns, indexes, views
- Almost every DBMS stores their catalog as a private database.
- Specialized code for “bootstrapping” catalog tables.

Data Representation

Data Representation

- A catalog table contain the schema information about the user tables
- The DBMS uses this schema information to figure out the tuple's **data representation**.
- In this way, it interprets the tuple's bytes into a set of attributes (types and values).

Variable Precision Numbers

- Inexact, variable-precision numeric type that uses the "native" C/C++ types.
 - ▶ Examples: FLOAT, REAL/DOUBLE
- Store directly as specified by **IEEE-754**.
- Typically faster than arbitrary precision numbers but can have rounding errors...

Variable Precision Numbers

Rounding Example

```
include <stdio.h>
```

```
int main(int argc, char* argv[]) {  
float x = 0.1;  
float y = 0.2;  
printf("x+y = %f\n", x+y);  
printf("0.3 = %f\n", 0.3);  
}
```

Output

```
x+y = 0.300000  
0.3 = 0.300000
```

Variable Precision Numbers

Rounding Example

```
include <stdio.h>
```

```
int main(int argc, char* argv[]) {  
float x = 0.1;  
float y = 0.2;  
printf("x+y = %.20f\n", x+y);  
printf("0.3 = %.20f\n", 0.3);  
}
```

Output

```
x+y = 0.30000001192092895508  
0.3 = 0.29999999999999998890
```


Postgres: Numeric

```
typedef unsigned char NumericDigit;
typedef struct {
    int ndigits;           // number of digits
    int weight;           // weight of 1st Digit
    int scale;            // scale factor
    int sign;             // positive/negative/NaN
    NumericDigit *digits; // digit storage
} numeric;
```


Schema Changes: Index

- CREATE INDEX:
 - ▶ Scan the entire table and populate the index (e.g., hash table: tuple id → tuple pointer).
 - ▶ Have to record changes made by txns that modified the table while another txn was building the index.
 - ▶ When the scan completes, lock the table and resolve changes that were missed after the scan started.
- DROP INDEX:
 - ▶ Just drop the index logically from the catalog.
 - ▶ It only becomes "invisible" when the txn that dropped it **commits**.
 - ▶ All ongoing txns will still have to update it.

Observation

- The relational model does **not** specify that we have to store all of a tuple's attributes together in a single page.
- This may **not** actually be the best **storage layout** for some workloads...

Workload Types

- OLTP: On-line Transaction Processing; Simple + Write-intensive
- OLAP: On-line Analytical Processing; Complex + Read-intensive
- HTAP: Hybrid Transaction/Analytical Processing; Medium + Mixed

Data Storage Models

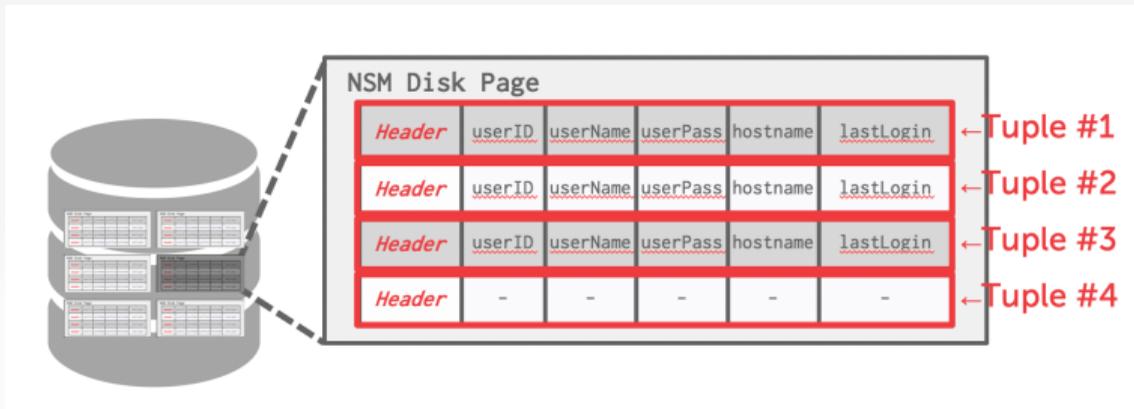
- The DBMS can store tuples in different ways that are better for either OLTP or OLAP workloads.
- We have been assuming the **n-ary storage model (NSM)** (*a.k.a.*, row storage) so far this semester.

N-ary Storage Model (NSM)

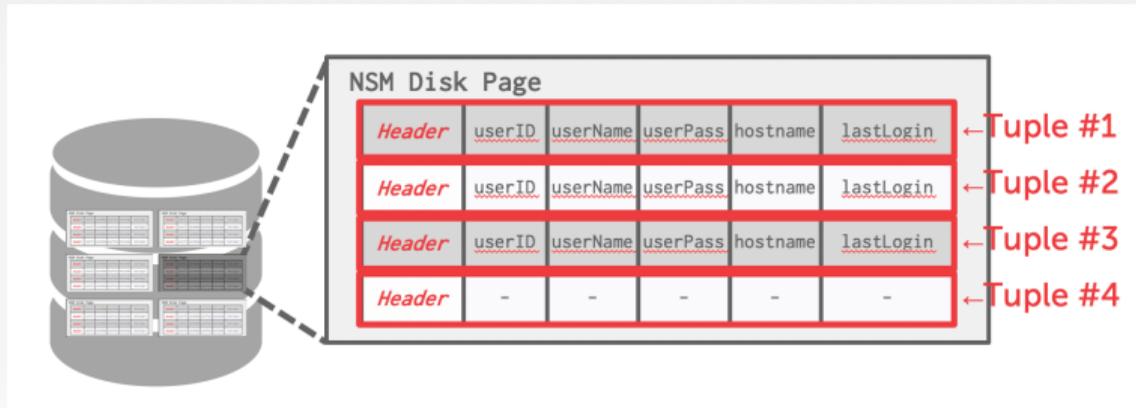
- The DBMS stores all attributes for a single tuple contiguously in a page.
- Ideal for OLTP workloads where queries tend to operate only on an individual entity and insert-heavy workloads.

N-ary Storage Model (NSM)

- The DBMS stores all attributes for a single tuple contiguously in a page.



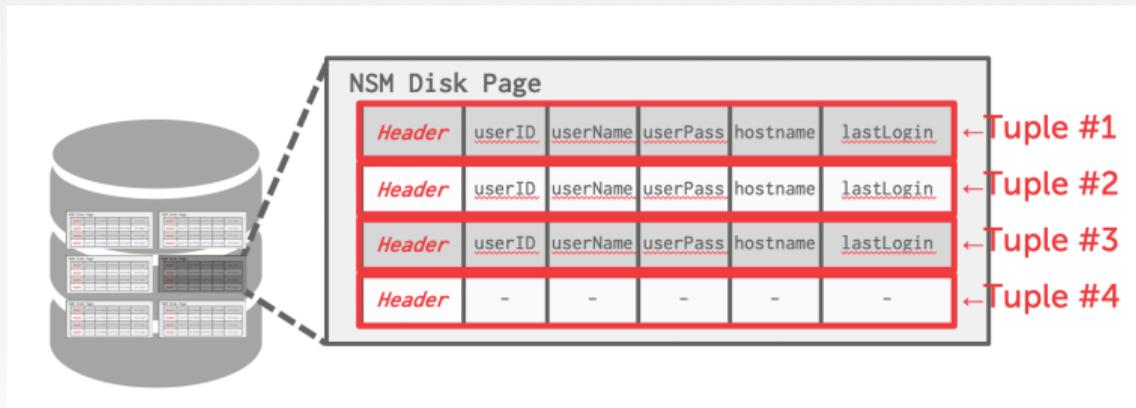
N-ary Storage Model (NSM)



```
SELECT * FROM useracct  
WHERE userName = ? AND userPass = ?
```

Use index to access the particular user's tuple.

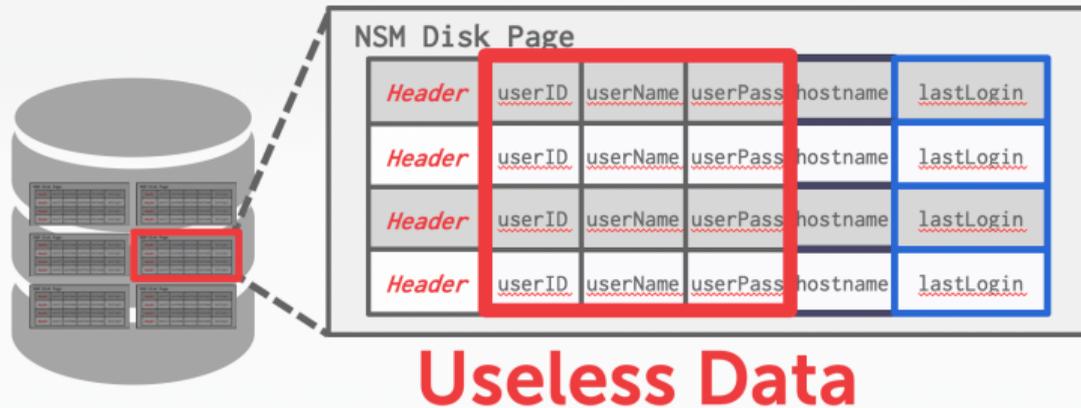
N-ary Storage Model (NSM)



`INSERT INTO useracct VALUES (?,?,...?)`

Add the user's tuple using `std::memcpy`.

N-ary Storage Model (NSM)



```
SELECT COUNT(U.lastLogin), EXTRACT(month FROM U.lastLogin) AS month  
FROM useracct AS U  
WHERE U.hostname LIKE '%.gov'  
GROUP BY EXTRACT(month FROM U.lastLogin)
```

Useless data accessed for this query.

N-ary Storage Model (NSM)

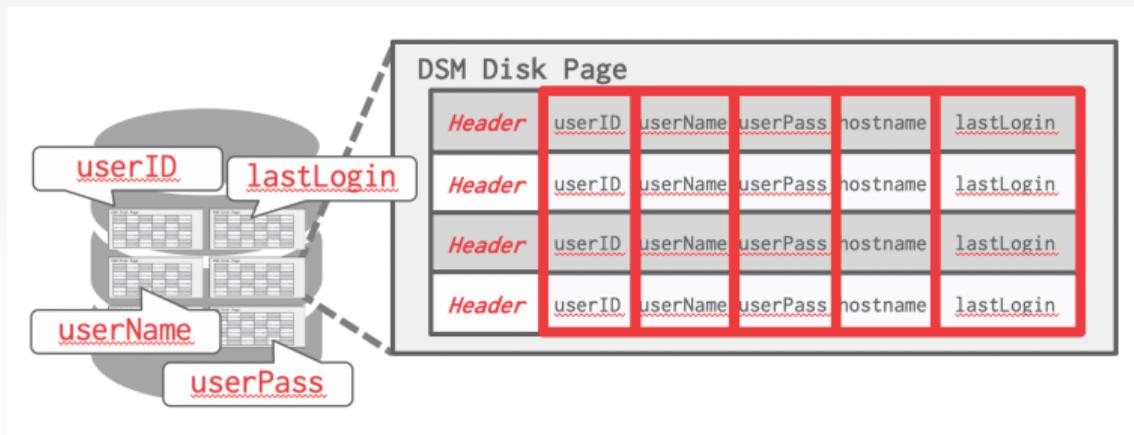
- Advantages
 - ▶ Fast inserts, updates, and deletes.
 - ▶ Good for queries that need the entire tuple.
- Disadvantages
 - ▶ Not good for scanning large portions of the table and/or a subset of the attributes.

Decomposition Storage Model (DSM)

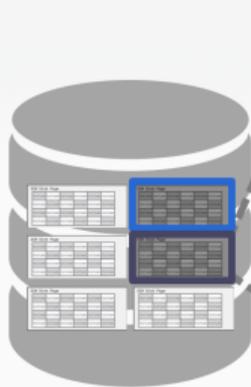
- The DBMS stores the values of a single attribute for all tuples contiguously in a page.
 - ▶ Also known as a "column store".
- Ideal for OLAP workloads where read-only queries perform large scans over a subset of the table's attributes.

Decomposition Storage Model (DSM)

- The DBMS stores the values of a single attribute for all tuples contiguously in a page.
 - ▶ Also known as a "column store".



Decomposition Storage Model (DSM)



DSM Disk Page

hostname	hostname	hostname	hostname	hostname	hostname
hostname	hostname	hostname	hostname	hostname	hostname
hostname	hostname	hostname	hostname	hostname	hostname
hostname	hostname	hostname	hostname	hostname	hostname

```
SELECT COUNT(U.lastLogin), EXTRACT(month FROM U.lastLogin) AS month  
FROM useracct AS U  
WHERE U.hostname LIKE '%.gov'  
GROUP BY EXTRACT(month FROM U.lastLogin)
```

Tuple Identification

- Choice 1: Fixed-length Offsets
 - ▶ Each value is the same length for an attribute.
- Choice 2: Embedded Tuple Ids
 - ▶ Each value is stored with its tuple id in a column.

Offsets

	A	B	C	D
0				
1				
2				
3				

Embedded Ids

	A	B	C	D
0		0		0
1		1		1
2		2		2
3		3		3

Decomposition Storage Model (DSM)

- Advantages
 - ▶ Reduces the amount wasted I/O because the DBMS only reads the data that it needs.
 - ▶ Better query processing and data compression (more on this later).
- Disadvantages
 - ▶ Slow for point queries, inserts, updates, and deletes because of tuple splitting/stitching.

DSM History

- 1970s: Cantor DBMS
- 1980s: **DSM Proposal**
- 1990s: SybaseIQ (in-memory only)
- 2000s: Vertica, VectorWise, MonetDB
- 2010s: Everyone

Next Class

- Buffer Management