The Legend of Drunken Query Master The Apprentice's Journey

Jay Pipes Community Relations Manager MySQL jay@mysql.com http://jpipes.com



These slides released under the Creative Commons Attribution-Noncommercial-Share Alike 3.0 License



- Who am I?
 - Just some dude who works at MySQL (eh...Sun)
 - Oh, I co-wrote a book on MySQL
 - Active PHP/MySQL community member
 - Other than that, semi-normal geek, married, 2 dogs, 2 cats, blah blah
- This talk is about understanding and tuning SQL applications





- Got a quick question?
 Just ask it.
- Got a longer question?



- Wait until the break or when the speaker forgets what he is talking about.
- Pictures of cute and cuddly animals are used frequently throughout presentation as a blatant attempt to improve your review of the speaker
 - If you hate cute and cuddly animals, you may want to leave now.



"...thou shalt not fight with the database, for it is thy friend"

can't we all just get along?



http://www.ringelkater.de/ringel_witzig/cat_dog_mouse.jpg

your friend, the database...

- Recognize the strengths and also the weaknesses of your database
- No database is perfect -- deal with it, you're not perfect either
- Think of both big things *and* small things
 - BIG: Architecture, surrounding servers, caching
 - SMALL: SQL coding, join rewrites, server config



- Some big things produce small results
 - Hardware changes can produce smaller improvements than many expect
- Some small things produce humungonormous results
 - Rewriting a poorly-constructed query can improve performance or scalability more than you might expect







keys to MySQL system architecture

- Understand storage engine abilities and weaknesses
- Understand how the query cache and important buffers works
- Understand optimizer's limitations
- Understand what should and should not be done at the application level
- If you understand the above, you'll start to see the database as a friend and not an enemy



"...poor skills in schema-kido shall lead to crushing defeat"

http://www.cold-moon.com/images/Motivators/GMs/defeat.jpg



DEFEAT

Sometimes you just should have seen it coming.



the schema

- Basic foundation of performance
- Everything else depends on it
- Choose your data types wisely
- Conquer the schema through partitioning



The Leaning Tower of Pisa *from Wikipedia*:

"Although intended to stand vertically, the tower began leaning to the southeast soon after the onset of construction in 1173 *due to a poorly laid foundation* and loose substrate that has allowed the foundation to shift direction."

smaller, smaller, smaller



The Pygmy Marmoset world's smallest monkey

This picture is a cheap stunt intended to induce kind feelings for the presenter.

Oh, and I *totally* want one of these guys for a pet.

The more records you can fit into a single page of memory/disk, the faster your seeks and scans will be

- Do you *really* need that **BIGINT**?
- Use **INT UNSIGNED** for IPv4 addresses
- Use VARCHAR carefully
 - Converted to CHAR when used in a temporary table
- Use **TEXT** sparingly
 - Consider separate tables
- Use **BLOB**s very sparingly
 - Use the filesystem for what it was intended

handling IPv4 addresses





- Drunken query master is not a big fan of SET and ENUM
- Sign of poor schema design
- Changing the definition of either will require a *full* rebuild of the table
- Search functions like FIND_IN_SET() are inefficient compared to index operations on a join



The Mandelbrot Set Wikipedia

Rumour has it that the Mandelbrot Set will be a full-fledged column type in MySQL 9.1, making for some very interesting application uses in database graphics processing...



how much storage space is consumed?

• With this definition, how many bytes will the "a" column consume per row?

```
CREATE TABLE t1 (
    a INT(1) UNSIGNED NOT NULL
);
```

- The number in parentheses is the ZEROFILL argument, not the storage space
- INT takes 4 bytes of space

```
- Regardless of the UNSIGNED or NOT NULL
```



"...table-kee-do shall show thee the Path of (De)Normalization"

fifth-degree black-belt in join-fu -->



Edgar F. Codd from Wikipedia:

"...while working for IBM, invented the relational model for database management, the theoretical basis for relational databases."



taking normalization way too far



http://thedailywtf.com/forums/thread/75982.aspx



divide et impera

- Vertical partitioning
 - Split tables with many columns into multiple tables
- Horizontal partitioning
 - Split table with many rows into multiple tables
- Partitioning in MySQL 5.1 is transparent horizontal partitioning within the DB...





Niccolò Machiavelli The Art of War, (1519-1520):

"A Captain ought, among all the other actions of his, endeavor *with every art to divide the forces of the enemy*, either by making him suspicious of his men in whom he trusted, or by giving him cause that he has to separate his forces, *and*, *because of this, become weaker*."

vertical partitioning



- Mixing frequently and infrequently accessed attributes in a single table?
- Space in buffer pool at a premium?
 - Splitting the table allows main records to consume the buffer pages without the extra data taking up space in memory
- Need **FULLTEXT** on your text columns?

09/18/08

zendcon08 - legend of drunken query master



http://obsidianwings.blogs.com/obsidian_wings/kitten.jpg

Think kittens get angry? Wait until you see what the query cache can do. "...thou shalt employ table-keedo in order to avoid the Fury of the Query Cache"

the MySQL query cache



- You must understand your application's read/write patterns
- Internal query cache design is a compromise between CPU usage and read performance
- Stores the MYSQL_RESULT of a SELECT along with a hash of the SELECT SQL statement
- Any modification to any table involved in the SELECT invalidates the stored result
- Write applications to be aware of the query cache
 - Use SELECT SQL_NO_CACHE

vertical partitioning ... continued

CREATE TABLE Products (product id INT NOT NULL **CREATE TABLE** Products (name VARCHAR(80) NOT NULL product id INT NOT NULL unit cost DECIMAL(7,2) NOT NULL , name VARCHAR(80) NOT NULL description **TEXT NULL** unit cost DECIMAL(7,2) NOT NULL image path **TEXT NULL** description TEXT NULL **PRIMARY KEY** (product id) image path TEXT NULL **INDEX** (name(20)) num views INT UNSIGNED NOT NULL num in stock INT UNSIGNED NOT NULL **CREATE TABLE** ProductCounts (num on order INT UNSIGNED NOT NULL product id INT NOT NULL **PRIMARY KEY** (product id) num views INT UNSIGNED NOT NULL **INDEX** (name(20)) num in stock INT UNSIGNED NOT NULL ENGINE=InnoDB; num on order INT UNSIGNED NOT NULL **PRIMARY KEY** (product id) // Getting a simple COUNT of products ENGINE=InnoDB: // easy on MyISAM, terrible on InnoDB **CREATE TABLE** TableCounts (**SELECT COUNT(*)** total products INT UNSIGNED NOT NULL **FROM** Products: ENGINE=MEMORY;

- Mixing static attributes with frequently updated fields in a single table?
 - Thrashing occurs with query cache. Each time an update occurs on any record in the table, all queries referencing the table are invalidated in the query cache
- Doing **COUNT(*)** with no **WHERE** on an indexed field on an InnoDB table?
 - Complications with versioning make full table counts very slow



"...thou shalt not be afraid of SQL, for it is thy strongest weapon"





coding like a join-fu master



Did you know? from *Wikipedia:*

Join-fu is a close cousin to **Jun Fan Gung Fu**, the method of martial arts Bruce Lee began teaching in 1959.

OK, not really.

- Be consistent (for crying out loud)
- Use ANSI SQL coding style
- Stop thinking in terms of iterators, for loops, while loops, etc
- Instead, think in terms of sets
- Break complex SQL statements (or business requests) into smaller, manageable chunks

SQL coding consistency

- Tabs and spacing
- Upper and lower case

Nothing pisses off the query master like inconsistent SQL code!

- Keywords, function names
- Some columns aliased, some not

```
SELECT
   a.first_name, a.last_name, COUNT(*) as num_rentals
FROM actor a
   INNER JOIN film f
    ON a.actor_id = fa.actor_id
GROUP BY a.actor_id
ORDER BY num_rentals DESC, a.last_name, a.first_name
LIMIT 10;
```

vs.

```
select first_name, a.last_name,
count(*) AS num_rentals
FROM actor a join film f on a.actor_id = fa.actor_id
group by a.actor_id order by
num_rentals DESC, a.last_name, a.first_name
LIMIT 10;
```

- Consider your teammates
- Like your programming code, SQL is meant to be read, not written

join-fu guidelines



See, even bears practice join-fu.

- Always try variations on a theme
- Beware of join hints
 - Can get "out of date"
- Just because it *can* be done in a single SQL statement doesn't mean it should
- Always test and benchmark your solutions
 - Use http_load (simple and effective for web stuff)

ANSI vs. Theta SQL coding style





Implicitly declare JOIN conditions in the WHERE clause

SELECT

```
a.first_name, a.last_name, COUNT(*) as num_rentals
FROM actor a, film f, film_actor fa, inventory i, rental r
WHERE a.actor_id = fa.actor_id
AND fa.film_id = f.film_id
AND f.film_id = i.film_id
AND r.inventory_id = i.inventory_id
GROUP BY a.actor_id
ORDER BY num_rentals DESC, a.last_name, a.first_name
LIMIT 10;
```

why ANSI style's join-fu kicks Theta style's ass

- MySQL only supports the INNER and CROSS join for the Theta style
 - But, MySQL supports the INNER, CROSS, LEFT, RIGHT, and NATURAL joins of the ANSI style
 - Mixing and matching both styles can lead to hard-toread SQL code
- It is supremely easy to miss a join condition with Theta style
 - especially when joining many tables together
 - Leaving off a join condition in the WHERE clause will lead to a cartesian product (not a good thing!)



"....Without the strength of explain-jitsu, thou shall perish in the Meadow of Misunderstanding"

chased by the Evil Army of Correlated Subqueries through the Meadow of Misunderstanding -->



http://i240.photobucket.com/albums/ff188/catsncheese/normal_domokuns-kitten.jpg



- Provides the execution plan chosen by the MySQL optimizer for a specific SELECT statement
- Simply append the word EXPLAIN to the beginning of your SELECT statement
- Each row in output represents a set of information used in the SELECT
 - A real schema table
 - A virtual table (derived table) or temporary table
 - A subquery in SELECT or WHERE
 - A unioned set



- **select_type** type of "set" the data in this row contains
- table alias (or full table name if no alias) of the table or derived table from which the data in this set comes
- type "access strategy" used to grab the data in this set
- **possible_keys** keys available to optimizer for query
- keys keys chosen by the optimizer
- rows estimate of the number of rows in this set
- Extra information the optimizer chooses to give you
- **ref** shows the column used in join relations

Example EXPLAIN output





Example #1 - the const access type

EXPLAIN SELECT * FROM rental WHERE rental id = $13 \setminus G$ id: 1 select type: SIMPLE table: rental type: const possible keys: PRIMARY key: PRIMARY key len: 4 ref: const rows: 1 Extra: 1 row in set (0.00 sec)

Constants in the optimizer

- a field indexed with a *unique non-nullable* key
- The access strategy of *system* is related to const and refers to when a table with only a single row is referenced in the SELECT
- Can be propogated across joined columns



Example #2 - constant propogation

```
EXPLAIN SELECT r.*, c.first name, c.last name
FROM rental r INNER JOIN customer c
ON r.customer_id = c.customer_id WHERE r.rental_id = 13\G
id: 1
 select_type: SIMPLE
      table: r
       type: const
possible keys: PRIMARY, idx fk customer id
        kev: PRIMARY
     key len: 4
       ref: const
       rows: 1
      Extra:
id: 1
 select type: SIMPLE
      table: c
       type: const
possible keys: PRIMARY
        key: PRIMARY
     key len: 2
        ref: const /* Here is where the propogation occurs...*/
       rows: 1
      Extra:
2 rows in set (0.00 \text{ sec})
```

Example #3 - the range access type

SELECT * FROM rental WHERE rental date **BETWEEN '2005-06-14' AND** '2005-06-16'\G id: 1 select type: SIMPLE table: rental type: range possible keys: rental date key: rental date key len: 8 ref: NULL rows: 364 Extra: Using where 1 row in set (0.00 sec)

Considerations with range accesses

- Index must be available on the field operated upon by a range operator
- If too many records are estimated to be returned by the condition, the range optimization won't be used
 - index or full table scan will be used instead
- The indexed field must not be operated on by a function call! (Important for all indexing)
The scan vs. seek dilemma

- A seek operation, generally speaking, jumps into a random place -- either on disk or in memory -- to fetch the data needed.
 - Repeat for each piece of data needed from disk or memory
- A scan operation, on the other hand, will jump to the start of a chunk of data, and sequentially read data -- either from disk or from memory -- until the end of the chunk of data
- For large amounts of data, scan operations tend to be more efficient than multiple seek operations



Example #4 - Full table scan

EXPLAIN SELECT * FROM rental WHERE rental date **BETWEEN** '2005-06-14' AND '2005-06-21'\G id: 1 select type: SIMPLE table: rental type: ALL possible keys: rental date /* larger range forces scan choice */ key: NULL key len: NULL ref: NULL rows: 16298 Extra: Using where 1 row in set (0.00 sec)

Why full table scans pop up

- No WHERE condition (duh.)
- No index on any field in WHERE condition
- Poor selectivity on an indexed field
- Too many records meet WHERE condition
- < MySQL 5.0 and using OR in a WHERE clause
- Using SELECT * FROM



Example #5 - Full index scan

EXPLAIN SELECT rental_id, rental_date FROM rental\G

id: 1
select_type: SIMPLE
table: rental
type: index
possible_keys: NULL
key: rental_date
key_len: 13
ref: NULL
rows: 16325
Extra: Using index
1 row in set (0.00 sec)



Example #6 - eq_ref strategy

```
EXPLAIN SELECT r.*, c.first name, <u>c.last name</u>
FROM rental r INNER JOIN customer c ON r.customer_id = c.customer_id
WHERE r.rental date BETWEEN '2005-06-14' AND '2005-06-16'\G
id: 1
 select type: SIMPLE
      table: r
       type: range
possible keys: idx fk customer id, rental date
        key: rental date
     key len: 8
        ref: NULL
       rows: 364
      Extra: Using where
id: 1
 select type: SIMPLE
      table: c
       type: eq ref
possible keys: PRIMARY
        key: PRIMARY
     key len: 2
        ref: sakila.r.customer id
       rows: 1
      Extra:
2 rows in set (0.00 \text{ sec})
```

09/18/08 OSCON 2007 - Target Practice

When eq_ref pops up

- Joining two sets on a field where
 - One side has unique, non-nullable index
 - Other side has at least a non-nullable index
- In example #6, an eq_ref access strategy was chosen because a unique, non-nullable index is available on customer.customer_id and an index is available on the rental.customer_id field

Nested loops join algorithm

- For each record in outermost set
 - Fetch a record from the next set via a join column condition
 - Repeat until done with outermost set
- Main algorithm in optimizer
 - Main work in 5.1+ is in the area of subquery optimization and additional join algorithms like semi- and merge joins



Example #7 - ref strategy

```
EXPLAIN SELECT * FROM rental
WHERE rental id IN (10,11,12)
AND rental date = '2006-02-01' \G
id: 1
 select type: SIMPLE
      table: rental
       type: ref
possible keys: PRIMARY, rental date
        key: rental_date
     key len: 8
        ref: const
       rows: 1
      Extra: Using where
1 row in set (0.02 \text{ sec})
```

OR conditions and the index merge

- Index merge best thing to happen in optimizer for 5.0
- Allows optimizer to use more than one index to satisfy a join condition
 - Prior to MySQl 5.0, only one index
 - In case of OR conditions in a WHERE, MySQL <5.0 would use a *full table scan*



Example #8 - index_merge strategy

```
EXPLAIN SELECT * FROM rental
WHERE rental id IN (10,11,12)
OR rental date = '2006-02-01' \G
id: 1
 select type: SIMPLE
       table: rental
        type: index merge
possible keys: PRIMARY, rental date
        key: rental date, PRIMARY
     key len: 8,4
        ref: NULL
        rows: 4
       Extra: Using sort union(rental date, PRIMARY); Using where
1 row in set (0.02 \text{ sec})
```



tired? break time...



http://jimburgessdesign.com/comics/images/news_pics/passed_out_corgi.jpg



ok, now that you've had a quick snack...





drunken query master says...

"...thou shall befriend the Index, for it is a Master of Joinfu and will protect thee from the Ravages of the Table Scan"

best coffee table. ever. -->



http://technabob.com/blog/wp-content/uploads/2008/05/nes_coffee_table.jpg

indexes - your schema's phone book

- Speed up SELECTs, but slow down modifications
- Ensure indexes on columns used in WHERE, ON, GROUP BY clauses
- Always ensure JOIN conditions are indexed (and have identical data types)
- Be careful of the column order
- Look for covering indexes
 - Occurs when all fields in one table needed by a SELECT are available in an index record



The Yellow Pages from *Wikipedia*:

"The name and concept of "Yellow Pages" came about in 1883, when a printer in Cheyenne, Wyoming working on a regular telephone directory ran out of white paper and used yellow paper instead"

selectivity - the key to good, er...keys

- Selectivity
 - % of distinct values in a column
 - -S=d/n
 - Unique/primary always 1.0
- If column has a low selectivity
 - It may still be put in a multi-column index
 - As a prefix?
 - As a suffix?
 - Depends on the application

remove crappy or redundant indexes

SELECT

t.TABLE SCHEMA AS `db`, t.TABLE NAME AS `table`, s.INDEX NAME AS `index name` , s.COLUMN NAME AS `field name`, s.SEQ IN INDEX `seq in index`, s2.max columns AS `# cols` , s.CARDINALITY AS `card`, t.TABLE ROWS AS `est rows` , ROUND(((s.CARDINALITY / IFNULL(t.TABLE ROWS, 0.01)) * 100), 2) AS `sel %` FROM INFORMATION SCHEMA.STATISTICS s **INNER JOIN** INFORMATION SCHEMA. TABLES t **ON** s.TABLE SCHEMA = t.TABLE SCHEMA **AND** s.TABLE NAME = t.TABLE NAME **INNER JOIN** (SELECT TABLE SCHEMA, TABLE NAME, INDEX NAME, MAX(SEQ IN INDEX) AS max columns **FROM** INFORMATION SCHEMA.STATISTICS WHERE TABLE SCHEMA != 'mysql' GROUP BY TABLE SCHEMA, TABLE NAME, INDEX NAME) AS s2 ON s.TABLE SCHEMA = s2.TABLE SCHEMA AND s.TABLE NAME = s2.TABLE NAME AND s.INDEX NAME = s2.INDEX NAME WHERE t.TABLE SCHEMA != 'mysql' /* Filter out the mysql system DB */ AND t.TABLE_ROWS > 10 /* Only tables with some rows */ AND s.CARDINALITY IS NOT NULL /* Need at least one non-NULL value in the field */ AND (s.CARDINALITY / IFNULL(t.TABLE ROWS, 0.01)) < 1.00 /* unique indexes are perfect anyway */ ORDER BY `sel %`, s.TABLE SCHEMA, s.TABLE NAME /* DESC for best non-unique indexes */ **LIMIT** 10;

TABLE_SCHEMA	TABLE_NAME	INDEX_NAME	COLUMN_NAME	SEQ_IN_INDEX	COLS_IN_INDEX	CARD	ROWS	+ SEL %
<pre>worklog planetmysql planetmysql planetmysql sakila sakila worklog worklog sakila mysqlforge</pre>	<pre>amendments entries entries entries inventory rental tasks tasks payment mw_recentchanges</pre>	<pre>text categories categories categories idx_store_id_film_id idx_fk_staff_id title title idx_fk_staff_id rc_ip</pre>	<pre>text categories title content store_id staff_id title description staff_id rc_ip</pre>	1 1 2 3 1 1 1 2 1 1	1 3 3 2 1 2 2 1 1 1	1 1 1 1 3 1 1 6 2	33794 4171 4171 4673 16291 3567 3567 15422 996	0.00 0.02 0.02 0.02 0.02 0.02 0.02 0.03 0.03

http://forge.mysql.com/tools/tool.php?id=85

indexed columns and functions don't mix

• A fast *range* access strategy is chosen by the optimizer, and the index on title is used to *winnow* the query results down



 A slow full table scan (the ALL access strategy) is used because a function (LEFT) is operating on the title column



solving multiple issues in a SELECT query

SELECT * FROM Orders WHERE TO_DAYS(CURRENT_DATE()) - TO_DAYS(order_created) <= 7;</pre>

• First, we are operating on an indexed column (order_created) with a function - let's fix that:

SELECT * FROM Orders WHERE order_created >= CURRENT_DATE() - INTERVAL 7 DAY;

• Although we rewrote the WHERE expression to remove the operating function, we still have a non-deterministic function in the statement, which eliminates this query from being placed in the query cache - let's fix that:

SELECT * FROM Orders WHERE order_created >= '2008-01-11' - INTERVAL 7 DAY;

- We replaced the function with a constant (probably using our application programming language). However, we are specifying SELECT * instead of the actual fields we need from the table.
- What if there is a **TEXT** field in Orders called order_memo that we don't need to see? Well, having it included in the result means a larger result set which may not fit into the query cache and may force a disk-based temporary table

SELECT order_id, customer_id, order_total, order_created
FROM Orders WHERE order_created >= '2008-01-11' - INTERVAL 7 DAY;



drunken query master says...

"...join-fu is thy best defense against the Evil Army of Correlated Subqueries"

general in the evil army -->

zendcon08 - legend of drunken query master





set-wise problem solving

"Show the last payment information for each customer"

CREATE TABLE `payment` (`payment id` smallint(5) unsigned NOT NULL auto_increment, customer id` smallint(5) unsigned NOT NULL, `staff id` tinyint(3) unsigned NOT NULL, `rental id` int(11) default NULL, `amount` decimal(5,2) NOT NULL, `payment date` datetime NOT NULL, `last update` timestamp NOT NULL ... on update CURRENT TIMESTAMP, PRIMARY KEY (`payment id`), KEY `idx_fk_staff_id` (`staff_id`), KEY `idx fk customer id` (`customer id`), **KEY** `fk payment rental` (`rental id`), **CONSTRAINT** `fk_payment_rental` **FOREIGN KEY** (`rental_id`) **REFERENCES** `rental` (`rental id`), **CONSTRAINT** `fk_payment_customer` **FOREIGN KEY** (`customer id`) **REFERENCES** `customer` (`customer id`) , **CONSTRAINT** `fk_payment_staff` **FOREIGN KEY** (`staff_id`) **REFERENCES** `staff` (`staff id`) **ENGINE**=InnoDB **DEFAULT CHARSET**=utf8

http://forge.mysql.com/wiki/SakilaSampleDB

09/18/08 zendcon08 - legend of drunken query master

thinking in terms of *foreach* loops...

OK, *for each* customer, find the maximum date the payment was made and get that payment record(s)

```
mysql> EXPLAIN SELECT
   -> p.*
   -> FROM payment p
   > WHERE p.payment date =
   -> ( SELECT MAX(payment date)
   -> FROM payment
   -> WHERE customer id=p.customer id
   -> )\G
id: 1
 select type: PRIMARY
      table: p
       type: ALL
       rows: 16567
      Extra: Using where
id: 2
 select type: DEPENDENT SUBQUERY
      table: payment
       type: ref
possible keys: idx fk customer id
        key: idx fk customer id
     kev len: 2
        ref: sakila.p.customer id
       rows: 15
2 rows in set (0.00 sec)
```

• A *correlated* subquery in the WHERE clause is used

 It will be reexecuted for each row in the primary table (payment)

 Produces 623 rows in an average of 1.03s



what about adding an index?

Will adding an index on (customer_id, payment_date) make a difference?

mysql> EXPLAIN SELECT	mysql> EXPLAIN SELECT
-> p.*	-> p.*
-> FROM payment p	-> FROM payment p
<pre>-> WHERE p.payment_date =</pre>	<pre>> WHERE p.payment_date =</pre>
-> (SELECT MAX(payment date)	-> (SELECT MAX(payment date)
-> FROM payment	-> FROM payment
-> WHERE customer_id=p.customer_id	-> WHERE customer_id=p.customer_id
->)\G	->)\G
**************************************	**************************************
id: 1	id: 1
select_type: PRIMARY	select_type: PRIMARY
table: p	table: p
type: ALL	type: ALL
rows: 16567	rows: 15485
Extra: Using where	Extra: Using where
**************************************	**************************************
id: 2	id: 2
<pre>select_type: DEPENDENT SUBQUERY</pre>	<pre>select_type: DEPENDENT SUBQUERY</pre>
table: payment	table: payment
type: ref	type: ref
<pre>possible_keys: idx_fk_customer_id</pre>	<pre>possible_keys: idx_fk_customer_id,ix_customer_paydate</pre>
<pre>key: idx_fk_customer_id</pre>	key: ix_customer_paydate
key_len: 2	key_len: 2
ref: sakila.p.customer_id	ref: sakila.p.customer_id
rows: 15	rows: 14
	Extra: Using index
2 rows in set (0.00 sec)	2 rows in set (0.00 sec)

Produces 623 rows in an average of 1.03s

09/18/08

zendcon08 - legend of drunken query master

Produces 623 rows in

an average of 0.45s

thinking in terms of sets...

OK, I have one set of last payments dates and another set containing payment information (so, how do I join these sets?)

mysgl> EXPLAIN SELECT -> p.* -> FROM (SELECT customer id, MAX(payment date) as last order -> FROM payment -> GROUP BY customer id ->) AS last orders -> INNER JOIN payment p -> ON p.customer id = last orders.customer id -> AND p.payment date = last orders.last order\G id: 1 select type: PRIMARY table: <derived2> type: ALL rows: 599 id: 1 select type: PRIMARY table: p type: ref possible keys: idx fk customer id, ix customer paydate key: ix customer paydate key len: 10 ref: last orders.customer id,last orders.last order rows: 1 id: 2 select type: DERIVED table: payment type: range key: ix_customer_paydate key len: 2 rows: 1107 Extra: Using index for group-by 3 rows in set (0.02 sec)

- A derived table, or subquery in the FROM clause, is used
- The derived table represents a set: last payment dates of customers
- Produces 623 rows in an average of 0.03s



drunken query master says...

"...join-fu shall assist you in your N:M relationships"

...but it won't help your other relationships



http://onlineportaldating.com/wp-content/uploads/2007/10/leaving.bmp

working with "mapping" or N:M tables



- The next few slides will walk through examples of querying across the above relationship
 - dealing with OR conditions
 - dealing with AND conditions



rows in

dealing with OR conditions

Grab all project names which are tagged with "mysql" OR "php"

<pre>mysql> SELECT p.n -> INNER JOIN -> ON p.proje -> INNER JOIN -> ON t2p.tag -> WHERE t.ta +</pre>	ame FROM Projec Tag2Project t2 ct_id = t2p.pro Tag t = t.tag_id g_text IN ('mys	ct p 2p oject sql','php');		
name + phpMyAdmin				
MySQL Stored Pr	ocedures Auto (Generator		
90 TOWS IN SEC (0	.05 Sec)			
++ id select_type ++	table type	+ possible_keys	 key	+ key_ +
	t I range	PRIMARY wix tag text	uiv tag text	1 52

ect_type	table	туре	possible_keys	кеу	key_len	ret	rows	Extra	
1PLE 1PLE 1PLE	t t2p p	r ange ref eq_ref	PRIMARY,uix_tag_text PRIMARY,rv_primary PRIMARY	uix_tag_text PRIMARY PRIMARY	52 4 4	NULL t.tag_id t2p.project	2 10 1	Using where Using index	
set (0.00	sec)								

 Note the order in which the optimizer chose to join the tables is exactly the opposite of how we wrote our SELECT



Grab all project names which are tagged with "storage engine" *AND* "plugin"

- A little more complex, let's grab the project names which match both the "mysql" tag and the "php" tag
- Here is perhaps the most common method - using a HAVING COUNT(*) against a GROUP BY on the relationship table
- EXPLAIN on next page...



the dang filesort

- The EXPLAIN plan shows the execution plan using a derived table containing the project IDs having records in the Tag2Project table with both the "plugin" and "storage engine" tags
- Note that a filesort is needed on the Tag table rows since we use the index on tag_text but need a sorted list of tag_id values to use in the join

id: 1 select type: PRIMARY table: <derived2> type: ALL rows: 2 id: 1 select type: PRIMARY table: p type: eq ref possible keys: PRIMARY key: PRIMARY key len: 4 ref: projects having all tags.project rows: 1 id: 2 select type: DERIVED table: t type: range possible keys: PRIMARY, uix tag text key: uix tag text key len: 52 rows: 2 Extra: Using where; Using index; Using temporary; Using filesort id: 2 select type: DERIVED table: t2p type: ref possible keys: PRIMARY key: PRIMARY key len: 4 ref: mysglforge.t.tag id rows: 1 Extra: Using index 4 rows in set (0.00 sec)



removing the filesort using CROSS JOIN

- We can use a CROSS JOIN technique to remove the filesort
 - We winnow down two copies of the Tag table (t1 and t2) by supplying constants in the WHERE condition
- This means no need for a sorted list of tag IDs since we already have the two tag IDs available from the CROSS JOINs...so no more filesort

mysq1 - - - - - - - - - - - - - - - - - - -	<pre>> EXPLAIN SELE > FROM Project > CROSS JOIN T > CROSS JOIN T > ON p.project > AND t2p.tag > INNER JOIN T > ON t2p.proje > AND t2p2.tag > WHERE t1.tag_t + AND t2.tag_t</pre>	CT p.name p ag t1 ag 2Projec _id = t2p = t1.tag ag2Projec ct = t2p2 = t2.tac _text = "si	e ct t2p p.project _id ct t2p2 2.project g_id "plugin" torage eng	jine";	+				++
id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1 1 1 1 +	SIMPLE SIMPLE SIMPLE SIMPLE SIMPLE SIMPLE	<pre>t1 t2 t2p t2p2 p sec)</pre>	const const ref eq_ref eq_ref	PRIMARY,uix_tag_text PRIMARY,uix_tag_text PRIMARY,rv_primary PRIMARY,rv_primary PRIMARY	uix_tag_text uix_tag_text PRIMARY PRIMARY PRIMARY	52 52 4 8 4	<pre>const const const const const,mysqlforge.t2p.project mysqlforge.t2p2.project</pre>	1 9 1 1	Using index Using index Using index Using index Using where

another technique for dealing with ANDs

 Do two separate queries - one which grabs tag_id values based on the tag text and another which does a self-join after the application has the tag_id values in memory

Benefit #1

 If we assume the Tag2Project table is updated 10X more than the Tag table is updated, the first query on Tag will be cached more effectively in the query cache

Benefit #2

 The EXPLAIN on the self-join query is *much* better than the HAVING COUNT(*) derived table solution



<pre>-> ON p.project_id = t2p.project</pre>
-> AND t2p.tag = 173
-> INNER JOIN Tag2Project t2p2
-> ON t2p.project = t2p2.project
-> AND t2p2.tag = 259;
++
name
++
Automatic data revision
memcache storage engine for MvSOL
++
2 rows in set (0.00 sec)

<pre>mysql> EXPLAIN SELECT p.name FROM Project p -> INNER JOIN Tag2Project t2p -> ON p.project_id = t2p.project -> AND t2p.tag = 173 -> INNER JOIN Tag2Project t2p2 -> ON t2p.project = t2p2.project -> AND t2p2.tag = 259;</pre>									
id	select_type	table	type	possible_keys	key	key_len	ref	rows	Extra
1 1 1	SIMPLE SIMPLE SIMPLE	t2p t2p2 p	ref eq_ref eq_ref	PRIMARY,rv_primary PRIMARY,rv_primary PRIMARY	PRIMARY PRIMARY PRIMARY	4 8 4	<pre>const const, mysqlforge.t2p.project mysqlforge.t2p2.project</pre>	9 1 1	Using index Using index Using where

understanding LEFT-join-fu



- Get the tag phrases which are not related to any project
- Get the tag phrases which are not related to any project OR the tag phrase is related to project #75
- Get the tag phrases not related to project #75

09/18/08 zendcon08 - legend of drunken query master

LEFT join-fu: starting simple...the NOT EXISTS

mysal> EXPLAIN SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.project IS NULL -> GROUP BY t.tag text\G id: 1 select type: SIMPLE table: t type: index possible keys: NULL key: uix tag text key len: 52 rows: 1126 Extra: Using index id: 1 select type: SIMPLE table: t2p type: ref possible keys: PRIMARY key: PRIMARY key len: 4 ref: mysqlforge.t.tag id rows: 1 Extra: Using where; Using index; Not exists 2 rows in set (0.00 sec) mysgl> SELECT -> t.tag_text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.project IS NULL -> GROUP BY t.tag text;

-> WHERE t2p.project IS NULL -> GROUP BY t.tag_text; +-----+ | tag_text | +----+ <snip> +-----+ 153 rows in set (0.01 sec)

- Get the tag phrases which are not related to any project
- LEFT JOIN ... WHERE x IS NULL
- WHERE x IS NOT NULL would yield tag phrases that *are* related to a project
 - But, then, you'd want to use an INNER JOIN

LEFT join-fu: a little harder

mysal> EXPLAIN SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.project IS NULL -> OR t2p.project = 75 -> GROUP BY t.tag text\G id: 1 select type: SIMPLE table: t type: index key: uix tag text key len: 52 ref: NULL rows: 1126 Extra: Using index id: 1 select type: SIMPLE table: t2p type: ref possible keys: PRIMARY key: PRIMARY key len: 4 ref: mysqlforge.t.tag id rows: 1 Extra: Using where; Using index 2 rows in set (0.00 sec) mysql> SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.project IS NULL -> OR t2p.project = 75 -> GROUP BY t.tag text; tag text <snip> 184 rows in set (0.00 sec)

 Get the tag phrases which are not related to any project OR the tag phrase is related to project #75

- No more NOT EXISTS optimization :(
- But, isn't this essentially a UNION?

LEFT join-fu: a UNION returns us to optimization

```
mysal> EXPLAIN SELECT
   -> t.tag text
   -> FROM Tag t
   -> LEFT JOIN Tag2Project t2p
   -> ON t.tag id = t2p.tag
   -> WHERE t2p.project IS NULL
   -> GROUP BY t.tag text
   -> UNION ALL
   -> SELECT
   -> t.tag text
   -> FROM Tag t
   -> INNER JOIN Tag2Project t2p
   -> ON t.tag id = t2p.tag
   -> WHERE t2p.project = 75 \setminus G
id: 1
 select type: PRIMARY
      table: t
       type: index
        key: uix tag text
     key len: 52
       rows: 1126
      Extra: Using index
  id: 1
 select type: PRIMARY
      table: t2p
       type: ref
        key: PRIMARY
     key len: 4
        ref: mysqlforge.t.tag id
       rows: 1
      Extra: Using where; Using index; Not exists
id: 2
 select type: UNION
      table: t2p
       type: ref
possible keys: PRIMARY,rv_primary
        key: rv primary
     key len: 4
        ref: const
       rows: 31
      Extra: Using index
```

```
id: 2
 select type: UNION
      table: t
      type: eq ref
possible keys: PRIMARY
       key: PRIMARY
    key len: 4
       ref: mysglforge.t2p.tag
      rows: 1
      Extra:
        id: NULL
 select type: UNION RESULT
      table: <union1,2>
5 rows in set (0.00 sec)
mysql> SELECT
   -> t.tag text
   -> FROM Tag t
   -> LEFT JOIN Tag2Project t2p
   -> ON t.tag id = t2p.tag
   -> WHERE t2p.project IS NULL
   -> GROUP BY t.tag text
   -> UNION ALL
   -> SELECT
   -> t.tag text
   -> FROM Tag t
   -> INNER JOIN Tag2Project t2p
   -> ON t.tag id = t2p.tag
   -> WHERE t2p.project = 75;
 tag text
+-----
<snip>
+---------+
184 rows in set (0.00 sec)
```

LEFT join-fu: the trickiest part...

mysal> SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.tag IS NULL -> AND t2p.project= 75 -> GROUP BY t.tag text; Empty set (0.00 sec) mysgl> EXPLAIN SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> WHERE t2p.tag IS NULL -> AND t2p.project= 75 -> GROUP BY t.tag text\G id: 1 select type: SIMPLE table: NULL type: NULL possible keys: NULL key: NULL key len: NULL ref: NULL rows: NULL Extra: Impossible WHERE noticed after reading const tables 1 row in set (0.00 sec)

- Get the tag phrases which are not related to project #75
- Shown to the left is the most common mistake made with LEFT JOINs
- The problem is where the filter on project_id is done...

LEFT join-fu: the trickiest part...fixed

mysal> EXPLAIN SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> AND t2p.project= 75 -> WHERE t2p.tag IS NULL -> GROUP BY t.tag text\G id: 1 select type: SIMPLE table: t type: index possible keys: NULL key: uix tag text key len: 52 rows: 1126 Extra: Using index id: 1 select type: SIMPLE table: t2p type: eq ref possible keys: PRIMARY, rv primary key: rv primary key len: 8 ref: const, mysqlforge.t.tag id rows: 1 Extra: Using where; Using index; Not exists 2 rows in set (0.00 sec) mysql> SELECT -> t.tag text -> FROM Tag t -> LEFT JOIN Tag2Project t2p -> ON t.tag id = t2p.tag -> AND t2p.project= 75 -> WHERE t2p.tag IS NULL -> GROUP BY t.tag text; tag text <snip> 674 rows in set (0.01 sec)

- Filters on the LEFT joined set must be placed in the ON clause
- Filter is applied before the LEFT JOIN and NOT EXISTs is evaluated, resulting in fewer rows in the evaluation, and better performance


Practical examples, but meant to show techniques of SQL problem solving

- Handling hierarchical queries
 - Adjacency lists
 - Nested sets
- Reporting query techniques
 - Running sums and aggregates
 - Ranking return results



drunken query master says...

"...join-fu and the Nested Sets Model shall shall deliver thee from the Adjacency List Model"

querying hierarchical structures

- Graphs and trees don't fit the relational model well
- Common solutions tend to use either of two techniques
 - Recursion (yuck.)
 - Application layer coding (ok.)
- A good solution blends two common tree-storage models
 - Adjacency list
 - Nested sets

adjacency list model

- Very common but doesn't scale
- Easy to query for:
 - Who is my parent?
 - Who are my children?
- Difficult to query for:
 - How many levels are in my tree?
 - Who are ALL the descendants of my grandfather's brother?

CR	EATE TABLE People (
	person_id INT UNSIGNED NOT NULL
,	name VARCHAR(50) NOT NULL
,	parent INT UNSIGNED NULL
,	PRIMARY KEY (person_id)
,	INDEX (parent)
)	ENGINE=InnoDB;

mysql> SELECT	<pre>* FROM People;</pre>	
person_id	name	parent
1 2 3 4 5 6 7	Great grandfather Grandfather Great Uncle Father Uncle Me Brother	NULL 1 2 2 4 4
7 rows in set	(0.00 sec)	++

adjacency list model - easy stuff

• Who is my parent?

• Who are my father's children?

 Who are my father's father's grandchildren?





adjacency list model - hard stuff

- How many levels in my hierarchy?
 - Told you. Yuck.
- Find all descendants of a specific person
 - Double yuck.
- Basic join-fu how not to do SQL?
 - Avoid cursors, iterators, etc

```
DELIMITER //
CREATE PROCEDURE get_max_levels()
BEGIN
SET @lowest_parent :=
  (SELECT MAX(parent) FROM People WHERE parent IS NOT NULL);
SET @levels := 1;
```

SET @current_parent = @lowest_parent;

```
WHILE @current_parent IS NOT NULL DO
   SET @current_parent :=
    (SELECT parent FROM People WHERE person_id = @current_parent);
   SET @levels := @levels + 1;
END WHILE;
```

```
SELECT @levels;
END //
```

END //

```
DELIMITER //
CREATE PROCEDURE get_node_descendants(IN to_find INT)
BEGIN
DROP TEMPORARY TABLE IF EXISTS child_ids;
CREATE TEMPORARY TABLE child_ids (child_id INT UNSIGNED NOT NULL);
...
WHILE @last_count_children > @new_count_children D0
...
INSERT INTO child_ids
SELECT person_id FROM new_children WHERE blah blah...;
SET @new_count_children := (SELECT COUNT(*) FROM child_ids);
END WHILE;
SELECT p.* FROM People
INNER JOIN child_ids
ON person_id = child_id;
```



nested sets model

- Uncommon because it is hard to grasp at first, but it really scales
- Easy to query for:
 - How many levels are in my tree?
 - Who are ALL the descendants of my grandfather's brother?
 - Various complex queries that would be impossible for the adjacency list model

CF	REATE TABLE People (
	person_id INT UNSIGNED NOT NULL
,	name VARCHAR(50) NOT NULL
,	left_side INT UNSIGNED NOT NULL
,	right_side INT UNSIGNED NOT NULL
,	PRIMARY KEY (person_id)
,	INDEX (parent)
)	ENGINE=InnoDB;

mysql> SELECT	<pre>* FROM People;</pre>	.		
person_id	name	parent		
1 2 3 4 5 6 7	Great grandfather Grandfather Great Uncle Father Uncle Me Brother	NULL 1 2 2 4 4		
7 rows in set (0.00 sec)				



- Each node in tree stores info about its location
 - Each node stores a "left" and a "right"
 - For the root node, "left" is always 1, "right" is always 2*n, where n is the number of nodes in the tree
 - For all other nodes, "right" is always equal to the "left" + (2*n) + 1, where n is the total number of child nodes of this node

- So...all "leaf" nodes in a tree have a "right" = "left" + 1

- Allows SQL to "walk" the tree's nodes
- OK, got all that? :)

nested sets model



- For the root node, "left" is always 1, "right" is always 2*n, where n is the number of nodes in the tree
- For all other nodes, "right" is always equal to the "left" + (2*n) + 1, where n is the total number of child nodes of this node



so, how is this easier?

- Easy to query for:
 - How many levels are in my tree?
 - Who are ALL the descendants of my grandfather's brother?
 - Various complex queries that would be impossible for the adjacency list model
- Efficient processing via set-based logic
 - Versus inefficient iterative/recursive model
- Basic operation is a BETWEEN predicate in a self join condition

nested list model - sets, not procedures...

- What is the depth of each node?
 - Notice the
 BETWEEN
 predicate in use
- What about the EXPLAIN output?
 - Oops
 - Add an index...

```
mysql> SELECT p1.person id, p1.name, COUNT(*) AS depth
    -> FROM People p1
   -> INNER JOIN People p2
   -> ON pl.left side BETWEEN p2.left side AND p2.right side
    -> GROUP BY p1.person id;
  person id |
           name
                             depth
           Great grandfather
         2 | Grandfather
                                 2
         3 | Great Uncle
                                 2
                                 3
         4 | Father
         5
                                 3
           Uncle
            Ме
                                 4
            Brother
id: 1
 select type: SIMPLE
      table: p1
       type: ALL
       rows: 7
      Extra: Using temporary; Using filesort
id: 1
 select type: SIMPLE
      table: p2
       type: ALL
       rows: 7
      Extra: Using where
ALTER TABLE People ADD UNIQUE INDEX ix nsm (left side, right side);
```



drunken query master says...

"...thou shalt build queries based on results you already know are correct"



- How do I find the max depth of the tree?
 - If the last query showed the depth of each node...then we build on the last query
- mysql> SELECT MAX(level) AS max_level FROM (
 -> SELECT pl.person_id, COUNT(*) AS level
 -> FROM People pl
 -> INNER JOIN People p2
 -> ON pl.left_side BETWEEN p2.left_side AND p2.right_side
 -> GROUP BY pl.person_id
 ->) AS derived;
 +-----+
 | max_level |
 +-----+
 | 4 |
 +-----+
 1 row in set (0.00 sec)
- Use this technique when solving set-based problems
 - Build on a known correct set and then intersect, union, aggregate, etc against that set

good, but could be better...

```
mysgl> EXPLAIN SELECT MAX(level) AS max level FROM (
   -> SELECT pl.person id, COUNT(*) AS level
   -> FROM People p1
   -> INNER JOIN People p2
   -> ON pl.left side BETWEEN p2.left side AND p2.right side
   -> GROUP BY pl.person id
   -> ) AS derived\G
id: 1
 select type: PRIMARY
      table: <derived2>
       type: ALL
       rows: 7
id: 2
 select type: DERIVED
      table: p1
      type: index
possible keys: ix nsm
       key: ix nsm
    key len: 8
       rows: 7
      Extra: Using index; Using temporary; Using filesort
id: 2
 select type: DERIVED
      table: p2
      type: index
possible keys: ix nsm
       key: ix nsm
    key len: 8
       rows: 7
      Extra: Using where; Using index
```

 Using covering indexes for everything

- "Using index"

 Unfortunately, we've got a filesort

- "Using filesort"

attacking unnecessary filesorts

mysgl> EXPLAIN SELECT MAX(level) AS max level FROM (-> SELECT pl.person id, COUNT(*) AS level -> FROM People p1 -> INNER JOIN People p2 -> ON pl.left side BETWEEN p2.left side AND p2.right side -> GROUP BY p1 person id -> ORDER BY NULL ->) AS derived\G id: 1 select type: PRIMARY table: <derived2> type: ALL rows: 7 id: 2 select type: DERIVED table: pl type: index possible keys: ix nsm key: ix nsm key len: 8 rows: 7 Extra: Using index; Using temporary; id: 2 select type: DERIVED table: p2 type: index possible keys: ix nsm key: ix nsm key len: 8 rows: 7 Extra: Using where; Using index

 GROUP BY implicitly orders the results

 If you don't need that sort, remove it it using ORDER BY NULL

finding a node's descendants

- Who are ALL my grandfather's descendants?
 - Remember the nasty recursive solution we had?

```
mysql> SELECT pl.name
    -> FROM People pl
    -> INNER JOIN People p2
    -> ON pl.left_side
    -> BETWEEN p2.left_side AND p2.right_side
    -> WHERE p2.person_id = @to_find
    -> AND pl.person_id <> @to_find;
+----+
| name
+----+
| Father
| Uncle
| Me
| Brother
+----+
4 rows in set (0.00 sec)
```

```
mysgl> EXPLAIN SELECT p1.name
   -> FROM People p1
   -> INNER JOIN People p2
   -> ON pl.left side BETWEEN p2.left side AND p2.right side
   -> WHERE p2.person id = @to find
   -> AND pl.person id <> @to find\G
      id: 1
 select type: SIMPLE
      table: p2
       type: const
possible keys: PRIMARY, ix nsm
        kev: PRIMARY
     key len: 4
        ref: const
       rows: 1
id: 1
 select type: SIMPLE
      table: p1
       type: range
possible keys: PRIMARY, ix nsm
        key: PRIMARY
     key len: 4
       rows: 4
      Extra: Using where
```

finding all nodes above a specific node

- Who are ALL my grandfather's predecessors?
- Look familiar to the last query?
 - What has changed?

mysql> -> -> -> -> -> ->	<pre>SELECT p2.name FROM People p1 INNER JOIN People p2 ON p1.left_side BETWEEN p2.left_side AND p2.right_side WHERE p1.person_id = @to_find AND p2.person_id <> @to_find;</pre>
name	
Grea	t grandfather
1 row :	in set (0.00 sec)

• What about now?

SELECT p2.name
FROM People p1
INNER JOIN People p2
ON p1.left_side
BETWEEN p2.left_side AND p2.right_side
WHERE p1.person_id = @to_find
AND p2.person_id <> @to_find;



- Lots more we could do with trees
 - How to insert/delete/move a node in the tree
 - How to connect the tree to aggregate reporting results
 - But not right now...
- Best practice
 - Use both adjacency list and nested sets for various query types
 - Little storage overhead
 - Best of both worlds



drunken query master says...

"...thou shalt study the practice of set-based formula replacement"



http://www.cat-pics.net/data/media/5/bottle%20drinking%20baby%20cat%20pics.jpg



formula replacement

- Take a formula you know works, and replace the variables with known sets
- Reduces errors significantly
- Forces you to think in terms of sets, instead of those darn FOR loops
- Examples:
 - Running aggregates
 - Ranking of results

reporting techniques

- Running aggregates
 - Without user variables
 - Running sums and averages
- Ranking of results
 - Using user variables
 - Using JOINs!



running aggregates

 When we want to have a column which "runs" a sum during the result set

SELECT MONTHNAME , COUNT(*) A FROM feeds WHERE create GROUP BY MON	(created) AS Added ed >= '20 NTH(creat	AS Month 007-01-01' ed);	????	+	
Month	Added		Month	Added	Total
January February March April May June	1 1 11 8 18 3	-	January February March April May June	1 11 8 18 3	1 2 13 21 39 42
6 rows in se	et (0.00	sec)	6 rows in s	et (0.00	sec)

basic formula for running aggregates

SELECT x1.key , x1.some_column , AGGREGATE_FN(x2.some_column) AS running_aggregate FROM x AS x1 INNER JOIN x AS x2 ON x1.key >= x2.key GROUP BY x1.key;

- Join a set (table) to itself using a >= predicate
 - ON x1.key >= x2.key
- Problem, though, when we are working with pre-aggregated data
 - Obviously, you can't do two GROUP BYs...

replacing sets in the running aggregate formula

SELECT

x1.key
, x1.some_column
, AGGREGATE_FN(x2.some_column)
FROM x AS x1
INNER JOIN x AS x2
ON x1.key >= x2.key
GROUP BY x1.key;

- Stick to the formula, but replace sets x1 and x2 with your preaggregated sets as derived tables
 - The right shows replacing x with derived

```
SELECT
FROM (
SELECT
  MONTH(created) AS MonthNo
, MONTHNAME(created) AS MonthName
  COUNT(*) AS Added
FROM feeds
WHERE created >= '2007-01-01'
GROUP BY MONTH(created)
) AS x1
INNER JOIN (
SELECT
  MONTH(created) AS MonthNo
, MONTHNAME(created) AS MonthName
  COUNT(*) AS Added
FROM feeds
WHERE created >= '2007-01-01'
GROUP BY MONTH(created)
) AS x2
ON \times 1.key >= \times 2.key
GROUP BY x1.key;
```

Finally, replace SELECT, ON and outer GROUP BY

Replace the greyed-out area with the correct fields

SELECT
x1.key
. x1.some column
AGGREGATE EN(x2.some_column)
FROM (
MUNIH(Created) AS MONTHNO
, MONTHNAME(created) AS MonthName
, COUNT(*) AS Added
FROM feeds
WHERE created >= $'2007-01-01'$
GROUP BY MONTH(created)
\land AC \vee 1
INNER JUIN (
SELECT
MONTH(created) AS MonthNo
, MONTHNAME(created) AS MonthName
. COUNT(*) AS Added
FROM feeds
WHERE created $> - '2007_01_01'$
(DOUD DV MONTU (areasted))
GRUUP BY MUNIH(Created)
) AS XZ
ON x1.key >= x2.key
GROUP BY x1.key;



and the running results...

+ MonthNo	HonthName	Added	++ RunningTotal			
1	January	1	1	-		
2	February	1	2			
3	March	11	13			
4	April	8	21			
5	May	18	39			
6	June	3	42			
++ 6 rows in set (0.00 sec)						

- Easy enough to add running averages
 Simply add a column for AVG(x2.Added)
- Lesson to learn: stick to a known formula, then replace formula elements with known sets of data (Keep it simple!)



ranking of results

- Using user variables
 - We set a @rank user variable and increment it for each returned result
- Very easy to do in both SQL and in your programming language code
 - But, in SQL, you can use that produced set to join with other results...

ranking with user variables

- Easy enough
 - But what about ties in the ranking?
- Notice that some of the films have identical prices, and so should be tied...
 - Go ahead and try to produce a *clean* way of dealing with ties using user variables...

mysql> SET @rank = 0; Query OK, 0 rows affected (0.00 sec)							
<pre>mysql> SELECT film_id, LEFT(title, 30) as title -> , rental_rate, (@rank:= @rank + 1) as rank -> FROM film -> ORDER BY rental_rate DESC -> LIMIT 10;</pre>							
film_id	title	rental_rate	rank				
243 DOORS PRESIDENT 7.77 1 93 BRANNIGAN SUNRISE 7.70 2 321 FLASH WARS 7.50 3 938 VELVET TERMINATOR 7.50 4 933 VAMPIRE WHALE 7.49 5 246 DOUBTFIRE LABYRINTH 7.45 6 253 DRIFTER COMMANDMENTS 7.44 7 676 PHILADELPHIA WIFE 7.44 8 961 WASH HEAVENLY 7.41 9 219 DEEP CRUSADE 7.40 10							
10 rows in set (0.00 sec) Hmm, I have to wonder what "Deep Crusade" is about							



- Again, we use a formula to compute ranked results
- Technique: use a known formulaic solution and replace formula values with known result sets
- The formula takes ties into account with the >= predicate in the join condition

```
SELECT
x1.key_field
, x1.other_field
, COUNT(*) AS rank
FROM x AS x1
INNER JOIN x AS x2
    ON x1.rank_field <= x2.rank_field
GROUP BY
x1.key_field
ORDER BY
x1.rank_field DESC;</pre>
```



replace variables in the formula

- SELECT
- x1.key_field
 , x1.other_field
 , COUNT(*) AS rank
 FROM x AS x1
 INNER JOIN x AS x2
 ON x1.rank_field <= x2.rank_field
 GROUP BY
 x1.key_field
 ORDER BY
 x1.rank_field DESCC
 LIMIT 10;</pre>



- Ties are now accounted for
- Easy to grab a "window" of the rankings
 - Just change LIMIT and OFFSET

SELECT
x1.film_id
, x1.title
, x1.rental_rate
, COUNT(*) AS rank
FROM film AS x1
INNER JOIN film AS x2
ON x1.rental_rate <= x2.rental_rate
GROUP BY
x1.film_id
ORDER BY
x1.rental_rate DESC
LIMIT 10;

+	+	+	+
film_id	title	rental_rate	rank
243 93 938 321 933 246 676 253 961 219	DOORS PRESIDENT BRANNIGAN SUNRISE VELVET TERMINATOR FLASH WARS VAMPIRE WHALE DOUBTFIRE LABYRINTH PHILADELPHIA WIFE DRIFTER COMMANDMENTS WASH HEAVENLY DEEP CRUSADE	7.77 7.70 7.50 7.50 7.49 7.45 7.44 7.44 7.41 7.41	1 2 4 4 5 6 8 8 9 10
+	+	+	+



refining the performance...

• EXPLAIN produces:

+ id	select_type	table	type	possible_keys	key	key_len	+ ref	rows	Extra
1	SIMPLE	x2	ALL	PRIMARY	NULL	NULL	NULL	952	Using temporary; Using filesort
1	SIMPLE	x1	ALL	PRIMARY	NULL		NULL	952	Using where

- And the query ran in 1.49s (that's bad, mkay...)
- No indexes being used
 - We add an index on film (film_id, rental_rate)

+ table	+ type	possible_keys	key	key_len	ref	rows	++ Extra
x2	index	ix_film_id	ix_film_id_rate	4	NULL	967	Using index; Using temporary; Using filesort
x1	ALL	ix_rate_film_id	NULL	NULL	NULL	967	Using where

- Results: slightly better performance of 0.80s
 - But, different GROUP and ORDER BY makes it slow



resources and thank you!

- PlanetMySQL
 - 300+ writers on MySQL topics
 - http://planetmysql.com
- MySQL Forge
 - Code snippets, project listings, wiki, worklog
 - http://forge.mysql.org



Baron Schwartz http://xaprb.com

MySQL performance guru and coauthor of High Performance MySQL, 2nd Edition (O'Reilly, 2008)

"xarpb" is Baron spelled on a Dvorak keyboard...