

Speed Issue with InnoDB

We created 16 identical InnoDB tables, here is some info for the first 4:

```
mysql> show create table i01\G
***** 1. row *****
      Table: i01
Create Table: CREATE TABLE `i01` (
  `objectId` bigint(20) NOT NULL AUTO_INCREMENT,
  `ra` float NOT NULL,
  `decl` float NOT NULL,
  `bMag` float NOT NULL,
  PRIMARY KEY (`objectId`),
  KEY `bMag` (`bMag`),
  KEY `ra` (`ra`)
) ENGINE=InnoDB AUTO_INCREMENT=30000001 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

mysql> show create table i02\G
***** 1. row *****
      Table: i02
Create Table: CREATE TABLE `i02` (
  `objectId` bigint(20) NOT NULL AUTO_INCREMENT,
  `ra` float NOT NULL,
  `decl` float NOT NULL,
  `bMag` float NOT NULL,
  PRIMARY KEY (`objectId`),
  KEY `bMag` (`bMag`),
  KEY `ra` (`ra`)
) ENGINE=InnoDB AUTO_INCREMENT=30000001 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

mysql> show create table i03\G
***** 1. row *****
      Table: i03
Create Table: CREATE TABLE `i03` (
  `objectId` bigint(20) NOT NULL AUTO_INCREMENT,
  `ra` float NOT NULL,
  `decl` float NOT NULL,
  `bMag` float NOT NULL,
  PRIMARY KEY (`objectId`),
  KEY `bMag` (`bMag`),
  KEY `ra` (`ra`)
) ENGINE=InnoDB AUTO_INCREMENT=30000001 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)

mysql> show create table i04\G
***** 1. row *****
      Table: i04
Create Table: CREATE TABLE `i04` (
  `objectId` bigint(20) NOT NULL AUTO_INCREMENT,
  `ra` float NOT NULL,
  `decl` float NOT NULL,
  `bMag` float NOT NULL,
  PRIMARY KEY (`objectId`),
  KEY `bMag` (`bMag`),
  KEY `ra` (`ra`)
) ENGINE=InnoDB AUTO_INCREMENT=30000001 DEFAULT CHARSET=latin1
1 row in set (0.00 sec)
```

We populated them with identical data, by loading the same csv file into each of them, here is a sample:

```
mysql> select count(*) from i01;
+-----+
| count(*) |
+-----+
| 30000000 |
+-----+
1 row in set (6.15 sec)

mysql> select count(*) from i02;
+-----+
| count(*) |
+-----+
| 30000000 |
+-----+
1 row in set (6.10 sec)

mysql> select count(*) from i03;
+-----+
| count(*) |
+-----+
| 30000000 |
+-----+
1 row in set (6.14 sec)

mysql> select count(*) from i04;
+-----+
| count(*) |
+-----+
| 30000000 |
+-----+
1 row in set (6.13 sec)

mysql> select * From i01 limit 10;
+-----+-----+-----+-----+
| objectId | ra | decl      | bMag |
+-----+-----+-----+-----+
| 1 | 0 | 0 | 0 |
| 2 | 0 | 89.097 | 20.47 |
| 3 | 0 | 64.7085 | 18.56 |
| 4 | 0 | 42.6267 | 20.29 |
| 5 | 0 | 36.4728 | 0 |
| 6 | 0 | 13.5169 | 0 |
| 7 | 0 | 12.9985 | 16.44 |
| 8 | 0 | -27.7431 | 0 |
| 9 | 0 | 66.5724 | 0 |
| 10 | 0 | 63.1921 | 17.62 |
+-----+-----+-----+-----+
10 rows in set (0.00 sec)

mysql> select * From i02 limit 10;
+-----+-----+-----+-----+
| objectId | ra | decl      | bMag |
+-----+-----+-----+-----+
| 1 | 0 | 0 | 0 |
| 2 | 0 | 89.097 | 20.47 |
| 3 | 0 | 64.7085 | 18.56 |
| 4 | 0 | 42.6267 | 20.29 |
| 5 | 0 | 36.4728 | 0 |
| 6 | 0 | 13.5169 | 0 |
| 7 | 0 | 12.9985 | 16.44 |
```

```

|      8 | 0 | -27.7431 | 0 |
|      9 | 0 | 66.5724 | 0 |
|     10 | 0 | 63.1921 | 17.62 |
+-----+-----+-----+-----+

```

10 rows in set (0.01 sec)

```
mysql> select * From i03 limit 10;
```

```

+-----+-----+-----+-----+
| objectId | ra | decl      | bMag |
+-----+-----+-----+-----+
|      1 | 0 | 0         | 0     |
|      2 | 0 | 89.097    | 20.47 |
|      3 | 0 | 64.7085   | 18.56 |
|      4 | 0 | 42.6267   | 20.29 |
|      5 | 0 | 36.4728   | 0     |
|      6 | 0 | 13.5169   | 0     |
|      7 | 0 | 12.9985   | 16.44 |
|      8 | 0 | -27.7431  | 0     |
|      9 | 0 | 66.5724   | 0     |
|     10 | 0 | 63.1921   | 17.62 |
+-----+-----+-----+-----+

```

10 rows in set (0.00 sec)

```
mysql> select * From i04 limit 10;
```

```

+-----+-----+-----+-----+
| objectId | ra | decl      | bMag |
+-----+-----+-----+-----+
|      1 | 0 | 0         | 0     |
|      2 | 0 | 89.097    | 20.47 |
|      3 | 0 | 64.7085   | 18.56 |
|      4 | 0 | 42.6267   | 20.29 |
|      5 | 0 | 36.4728   | 0     |
|      6 | 0 | 13.5169   | 0     |
|      7 | 0 | 12.9985   | 16.44 |
|      8 | 0 | -27.7431  | 0     |
|      9 | 0 | 66.5724   | 0     |
|     10 | 0 | 63.1921   | 17.62 |
+-----+-----+-----+-----+

```

10 rows in set (0.00 sec)

```
mysql>
```

```
mysql> select count(*) from i01 where bMag>0.2;
```

```

+-----+
| count(*) |
+-----+
| 13331273 |
+-----+

```

1 row in set (3.38 sec)

```
mysql> select count(*) from i02 where bMag>0.2;
```

```

+-----+
| count(*) |
+-----+
| 13331273 |
+-----+

```

1 row in set (3.38 sec)

```
mysql> select count(*) from i03 where bMag>0.2;
```

```

+-----+
| count(*) |
+-----+

```

```

+-----+
| 13331273 |
+-----+
1 row in set (3.43 sec)

mysql> select count(*) from i04 where bMag>0.2;
+-----+
| count(*) |
+-----+
| 13331273 |
+-----+
1 row in set (3.38 sec)

mysql> select count(*) from i01 where ra<3;
+-----+
| count(*) |
+-----+
| 4675829 |
+-----+
1 row in set (1.08 sec)

mysql> select count(*) from i02 where ra<3;
+-----+
| count(*) |
+-----+
| 4675829 |
+-----+
1 row in set (1.08 sec)

mysql> select count(*) from i03 where ra<3;
+-----+
| count(*) |
+-----+
| 4675829 |
+-----+
1 row in set (1.08 sec)

mysql> select count(*) from i04 where ra<3;
+-----+
| count(*) |
+-----+
| 4675829 |
+-----+
1 row in set (1.08 sec)

```

So what is the issue? For some reason optimizer will pick a different plan for different tables:

```

mysql> explain select count(*) from i01, i02 where i01.ra=i02.ra and i01.ra<3 and i02.bMag>0.2;
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i02 | ALL | bMag,ra | NULL | NULL | NULL | 30000370 | Using where |
| 1 | SIMPLE | i01 | ref | ra | ra | 4 | lsst_1.i02.ra | 17 | Using index |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.01 sec)

mysql> explain select count(*) from i02, i03 where i02.ra=i03.ra and i02.ra<3 and i03.bMag>0.2;

```

```

+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i03 | range | bMag,ra | bMag | 4 | NULL | 5006974 | Using where |
| 1 | SIMPLE | i02 | ref | ra | ra | 4 | lsst_1.i03.ra | 8 | Using index |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> explain select count(*) from i03, i04 where i03.ra=i04.ra and i03.ra<3 and i04.bMag>0.2;
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i04 | ALL | bMag,ra | NULL | NULL | NULL | 30000445 | Using where |
| 1 | SIMPLE | i03 | ref | ra | ra | 4 | lsst_1.i04.ra | 19 | Using index |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> explain select count(*) from i04, i05 where i04.ra=i05.ra and i04.ra<3 and i05.bMag>0.2;
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i05 | ALL | bMag,ra | NULL | NULL | NULL | 30000507 | Using where |
| 1 | SIMPLE | i04 | ref | ra | ra | 4 | lsst_1.i05.ra | 12 | Using index |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> explain select count(*) from i05, i06 where i05.ra=i06.ra and i05.ra<3 and i06.bMag>0.2;
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i06 | range | bMag,ra | bMag | 4 | NULL | 4611370 | Using where |
| 1 | SIMPLE | i05 | ref | ra | ra | 4 | lsst_1.i06.ra | 14 | Using index |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> explain select count(*) from i06, i07 where i06.ra=i07.ra and i06.ra<3 and i07.bMag>0.2;
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| id | select_type | table | type | possible_keys | key | key_len | ref | rows | Extra |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| 1 | SIMPLE | i07 | ALL | bMag,ra | NULL | NULL | NULL | 30000474 | Using where |
| 1 | SIMPLE | i06 | ref | ra | ra | 4 | lsst_1.i07.ra | 15 | Using index |
+----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
2 rows in set (0.00 sec)

```

And it might take 20 sec (the plan with "ALL"), or 17 min (the plan with "range").