Prof. Dr. Stefan Brass
Institut für Informatik
MLU Halle-Wittenberg

November 17, 2017

Databases IIB: DBMS-Implementation — Exercise Sheet 6 —

Part a) to e) will be discussed in class, you only have to submit Part f) and g). But please think about the questions in a) before the meeting. Send your solution for the homework to the instructor via EMail: `brass@informatik.uni-halle.de` (with "`dbi17`" in the subject line). The official deadline is November 22, 12:00.

Repetition Questions

a) What would you answer to the following questions in an oral exam?

- Name some background processes of Oracle and explain their purpose. In order to be not too Oracle-specific, think about tasks a DBMS should do asynchronously, i.e. not as part of the process or thread that executes a query or update.
- Name some data structures in the shared memory of all Oracle processes (the System Global Area). Again, try to think more generally about any DBMS that executes several queries in different processes/threads at the same time. What data should they share? Even if you have a server that executes only one query at a time, what data should it keep in memory from one query to the next?
- What is the purpose of a checkpoint?
- What is the difference between the dedicated server architecture and the multi-threaded server architecture of Oracle?
- What is the purpose of the archiver process in Oracle? Why should you switch a database with important data to the `ARCHIVELOG` mode? Against which failures is the data protected when not in `ARCHIVELOG` mode?
- What are the main components of a disk (hard disk drive)?
- What is a sector? How can one specify the position of a single sector on the disk, i.e. what are the components of a sector address?
- What is a typical sector size? Why do many DBMS and operating systems use a larger block size ("cluster size" on Windows)? What is a typical block size? Discuss advantages/disadvantages of increasing the block size.
- Why do some disks have more sectors on the outer tracks than on the inner tracks?
- If a power failure happens, is it possible that a block is only partially written?

In-Class Exercises

b) Check again the Oracle processes that are running on the server machine and find out, whether our installation uses the dedicated server architecture or the multi-threaded server architecture:

```
ps -ef | fgrep ora
```

c) Open three terminal windows on the virtual server machine. (Probably, it would be time to create a small shell script `vm` that contains the `ssh` command for logging into the server, then you do not have to remember the IP-number in future. Use "`chmod u+x vm`" to make the file executeable. Then enter "`./vm`" to execute it.)

- In one window, log into SQL*Plus as a normal user, and start a transaction, e.g. update a row. Make sure that you do not have `AUTOCOMMIT` mode (then SQL*Plus would send a `COMMIT` to the server after any update). E.g. you can enter (inside SQL*Plus):

  ```
  show autocommit
  ```

  If necessary, enter

  ```
  set autocommit off
  ```

- If you never tried what happens when you want to access a row that is locked by somebody else, this would be a good time. Log into SQL*Plus in the second window as the same user, and first write a query that shows the updated row. You will notice that it appears unchanged as long as you do not enter `COMMIT;` in the first window. Then enter an `UPDATE` statement for the same row. It will appear that the system "hangs", i.e. there is no reaction until you enter `COMMIT;` or `ROLLBACK;` in the first window and thereby release the lock. Then the update will proceed as normal. However, please create a situation where one session waits for a lock hold by the other session.
- In the third window, log in as administrator and use the tables `V$TRANSACTION`, `V$SESSION`, `V$TRANSACTION_ENQUEUE`, and `V$LOCK` to find out (1) who is currently logged in, (2) who has an active transaction, and (3) who holds locks that other users wait for.

d) How much memory is reserved for database buffers (the cache for datafile blocks)? You can get this information from `V$SGA` (or `V$SGASTAT`). You can also use the UNIX/Linux command "`top`" to see the amount of memory of the machine that is actually used (this should be larger, of course). Check the documentation for the parameters

- `DB_BLOCK_SIZE`:
  [https://docs.oracle.com/cd/B28359_01/server.111/b28320/initparams047.htm]
- `DB_BLOCK_BUFFERS`:
  [https://docs.oracle.com/cd/B28359_01/server.111/b28320/initparams044.htm]

- `DB_CACHE_SIZE`:
  [https://docs.oracle.com/cd/B28359_01/server.111/b28320/initparams049.htm]
- `SGA_TARGET`:
  [https://docs.oracle.com/cd/B28359_01/server.111/b28320/initparams222.htm]
- `MEMORY_TARGET`:
  [https://docs.oracle.com/cd/B28359_01/server.111/b28320/initparams133.htm]

You can check the values of these parameters with the "`show parameter`" command of SQL*Plus, or by directly selecting from `V$PARAMETER`. Note that it suffices to specify a substring of the parameter name with "`show parameter`". You will find that the database uses "Automatic Shared Memory Management":

[https://docs.oracle.com/cd/B28359_01/server.111/b28310/memory004.htm]

This is already not the fully automatic version that is available in Oracle 11g. That version includes also memory inside the server processes.

If you are really interested, and want to understand the `CACHE_SIZE` parameter, you should read this section about nonstandard block sizes:

[https://docs.oracle.com/cd/B28359_01/server.111/b28310/tspaces003.htm]

However, having different block sizes in one Oracle database is uncommon, and mainly intended for migration.

e) Log in as `SYSTEM` and grant your normal user account the role `PLUSTRACE`. Then log in as the normal user and enter

```
SET AUTOTRACE ON STATISTICS
```

Run a query. The output for "`physical reads`" shows the number of blocks actually read from disk. The output for "`db block gets`" and "`consistent gets`" shows the number of blocks accessed in total (using also blocks that are already cached in memory). The difference between both values is that "`db block gets`" are accesses to the current version of a block, whereas "`consistent gets`" are accesses to the version that contains only committed changes when the current query started. Note that also accesses to the data dictionary (including information needed by the query optimizer) count. If you immediately rerun the same query (using the cached execution plan), the numbers will be much smaller. You find more information in the manual:

[https://docs.oracle.com/cd/E18283_01/server.112/e16604/ch_eight.htm]

Homework Exercise

f) Please either find out which disk is installed on your own PC, or name some disk you would like in your next PC. Furthermore, try to find some data on that disk in the internet. Any three data items (numbers) would suffice (at least the disk capacity you should know, maybe you can find e.g. the seek time and the rotation speed in addition). If you cannot find data on your disk, choose any other disk.

On Windows, the name of the disk should be listed in the "Device Manager", which you can reach via "Settings", "System", "about", and then it is linked at the bottom. For Linux, you find some information on the following pages:

- [https://www.cyberciti.biz/faq/find-hard-disk-hardware-specs-on-linux/]
- [https://unix.stackexchange.com/questions/4561/]

Some places for getting performance comparisons of disks are:

- [http://www.storagereview.com/]
- [http://www.tomshardware.com/]

Of course, you can also check the web pages of big disk manufacturers for data sheets of their disks, e.g.

- [https://www.seagate.com/internal-hard-drives/hdd/] (English)
  [https://www.seagate.com/de/de/internal-hard-drives/hdd/] (German)
- [https://www.hgst.com/products/hard-drives] (English)
  [https://www.hgst.com/de/products/hard-drives] (German)
- [https://www.wdc.com/products/internal-storage.html] (English)
  [https://www.wdc.com/de-de/products/internal-storage.html] (German)
- [https://en.wikipedia.org/wiki/List_of_computer_hardware_manufacturers]

g) Read the first chapter of the "Oracle Database Concepts Manual"

[https://docs.oracle.com/cd/E11882_01/server.112/e40540/toc.htm]

and answer the following questions:

- In which version Oracle got "PL/SQL"?
- Compare the architecture diagram with the one on the slides and name one process or memory area that is missing on the slides.

Read also the first part of Chapter 8 about "Server-Side Programming" and do the following:

- Name at least one advantage of server-side programming (compared with programming the functionality in the application program that runs on the client).