THESIS Evaluating and Comparing Performance of Database Systems
STUDENT David Phung and Tobias Ronge
SUPERVISOR Per Andersson (LTH), Sverrir Valgeirsson (Consafe Logistics)
EXAMINER Flavius Gruian (LTH)

A Method for Performance Change Assessment Before a DBMS Upgrade

POPULAR SCIENCE CONCLUSION **David Phung and Tobias Ronge**

In this thesis, instructions for how to perform such impact analysis have been presented, including an application for performance measuring. Such an analysis would make the database developers and administrators feel safer that the database works as before, after an upgrade has been done. And while the area of such impact analysis is huge - it includes timing, memory allocation, output and syntax differences - the work of this thesis can be seen as a step forward in establishing methods for studying these impacts.

Databases are used in a lot of different industries today. If a database fails because of, say, logical errors or lack of memory, there can be devastating consequences. There are several known database management systems available today and many of them are frequently updated. Sometimes, even a switch to a completely different system might be wished for. However, such upgrades and switches are likely to have impacts on the functionality and performance which is why impact analysis should be considered before an upgrade is fully accepted.

In the thesis, a method based on capture and replay - that is, capturing commands sent to the database and replay them - is presented for performance testing, more precisely for measuring memory and timing. Using a test application programmed specifically for this thesis, performance testing can be done on different versions of a well known database management system. By first recording database communication in a so-called trace file, the communication can be replayed on the database a large number of times, resulting in a realistic simulation of real communication while making performance measurements during the execution. The application allows for multiple sequential runs and also presents statistics on the results.

Advantages of the method is that the tester can choose the test environment on which to run the commands, along with the fact that the set of commands can be run over and over again. Using the application, the developer does not need to temporarily take the original database system out of action in order to test it.

The application has been programmed with the intent of allowing future programmers to contribute to the code, allowing for more complex testing and more database management systems to be compatible.