Interactive Multi-Objective Query Optimization in Mobile-Cloud Database Environments based on a Weighted Sum Model

Florian Helff[1] Chenxiao Wang[1] Le Gruenwald[1] Laurent d'Orazio[2]

[1] School of Computer Science
University of Oklahoma
Norman, Oklahoma, USA
{fhelff, chenxiao, ggruenwald}@ou.edu

[2] CNRS, UMR 6158, LIMOS
Blaise Pascal University
Clermont-Ferrand, France
laurent.dorazio@isima.fr

ABSTRACT

Multiple cost objectives such as monetary cost, query execution time and mobile device energy consumption have to be considered for query optimization in mobile-cloud database environments where multiple users on mobile devices request services executed on a cloud. Requested data might be partially cached on the mobile device itself or has to be processed on the cloud which leads to those various costs. Choosing an optimal query execution plan is crucial to minimize the overall cost but is related to user preferences on those various costs. This paper presents a user-interactive multi-objective query optimization strategy based on a modified weighted sum model, called Normalized Weighted Sum Model (NWSM). The focus of this paper is the user interaction with the query optimization strategy and the comparison to the existing interactive multi-objective optimization approach, Skyline Queries. The evaluation of this analysis is supported by a user study comparing the accuracy of the user's decision and the amount of time the user needs to make such a decision using both approaches.

1. INTRODUCTION

An example architecture of a mobile-cloud database environment [1] is shown in Figure 1. In this environment, a user issues queries from a mobile device to obtain data. This data is either stored on the cloud or retrieved from a cache on the mobile device. In the following, the query optimizer generated many query execution plans (QEPs) for this query. Each QEP is associated with three different costs: the monetary cost for query execution on the cloud, the overall query execution time, and the energy consumption on the mobile device where the query might be executed [2].These three costs constitute the three cost objectives that the query optimizer needs to minimize in order to choose the optimal query execution plan among the others it generates. Considering different cloud pricing models [3], this optimization process is a stretch of contradicting objectives.

In an interactive query optimization system, it is the task of a user to interact with the system to find an optimal QEP with minimized costs according to his preferences on those three objectives.

A user-friendly decision strategy to find the optimal QEP with its cost is the lexicographical ordering [4]. This strategy focuses on a single main objective, such as execution time, and orders further objectives, like monetary cost and energy consumption, in a descending order. The user-interaction is simple since only the order of objectives has to be selected. Nevertheless using the lexicographical ordering is not sufficient to minimize the different costs as shown in following example:

Consider the three query execution plans (QEPs) with their costs

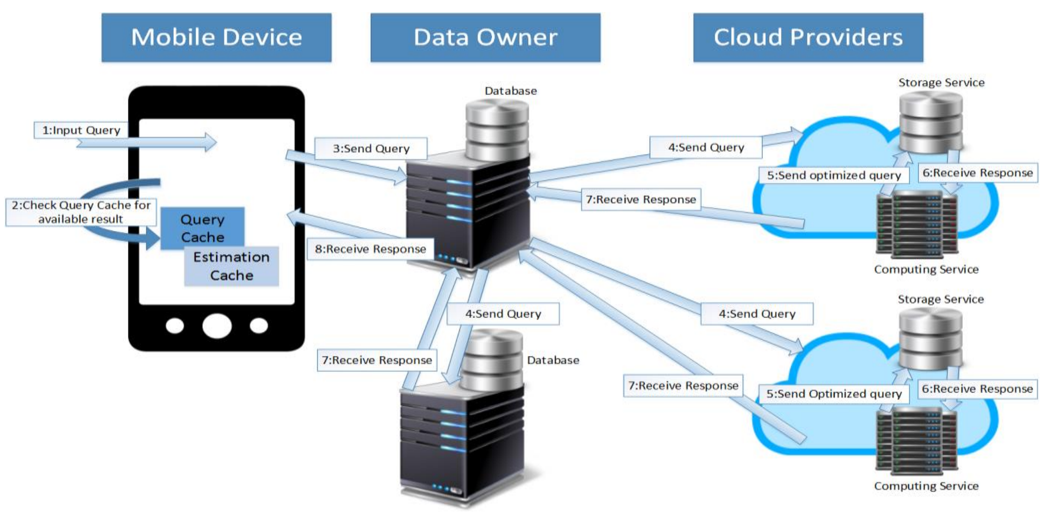

Figure 1: Mobile - Cloud Database Environment [1]

for monetary costs (M), execution time (T) and energy consumption (E) shown in Figure 2.

QEP1: {M= \$0.080; T= 0.5s; E= 0.012 mA}

QEP2: {M= \$0.050; T= 3.0s; E= 0.300 mA}

QEP3: {M= \$0.055; T= 0.6s; E= 0.013 mA}

Figure 2: Execution Plan Costs Example

Asking a user for his order on those three objectives will always results in a decision selecting either QEP1 or QEP2 for execution since those QEPs have a minimum cost in one of the three objectives. Nevertheless, QEP3 would be a competitive choice considering all three objectives equally. This shows that a one dimensional optimization strategy is easy to handle for a user but does not give the user the option to consider multiple objectives at the same time.

A different multi-objective decision approach is the Skyline Query strategy which considers all objectives at the same time. Eliminating dominated QEPs, a skyline query results in a set of QEPs which are optimal in its combination of different costs [5]. It is the task of the user to determine one QEP of this set to execute. This approach has two weaknesses. The first weakness is the size of the skyline. The size is dependent on the number of objectives and on the number of different QEPs which is related to the pricing models. The resulting decision for a user to select one QEP out of this set can be unsatisfying since this set tends to be large [6]. Approaches are made to reduce the size of the skyline which are discussed in section 2. Furthermore, it is a weakness of the skyline approach that users have to be aware of the overall cost constraints to choose accordingly. In the case of multiple users within an organization, users might not know the constraints for their needed query execution.

In summary, the skyline approach makes it possible for a user to consider multiple objectives at the same time but also burdens the user with the decision to select the final QEP.

In our previous work, we presented the Normalized Weighted Sum Algorithm (NWSA) [7]. This algorithm allows the user to make a decision for the QEP based on all objectives but does not burden the user with the task to select a QEP out of a set as seen in the skyline query approach. Using NWSA, the user assigns weights to each objective which are then multiplied with the different costs and accumulated result in a score for each QEP. The QEP with the lowest score will then be selected and executed. NWSA also allows a separation of different users: those who execute queries and those who set up weight profiles. Built on the promising results of NWSA, in this paper, we focus on the human-in-the-loop nature of multi-objective query optimization. Specifically, we focus on the comparison of the fundamental different user interactions with NWSA and with the general Skyline approach when considering multi-objective query optimization.

The rest of the paper is structured as follows. Section 2 covers the Pareto Set, Skyline Queries and their related work of multi-objective query optimization. Section 3 gives a brief overview over the predecessor research of the NWSA [7] and focusses on the comparison of user interaction of NWSA and Pareto Set approaches. Section 4 describes the evaluation of this analysis based on a user studies testing for accuracy and user interaction duration of both approaches. The conclusion and an outlook on future work are given in section 5.

2. PARETO SET AND SKYLINE QUERIES

Multi-Objective Optimization is a crucial problem to the mobile-cloud environment since multiple contradicting objectives have to be balanced.

The Pareto Set is one alternative to filter dominated alternatives out of possible solutions. An alternative 'A' is dominating an alternative 'B' if at least one objective (decision variable) of 'A' is better than the objective in 'B' and all other objectives in 'A' are at least equal to the objectives in 'B'. Those dominating alternatives are called Pareto optimal as defined by Zitzler and Thile [8].

When using the Pareto Set approach for query optimization, every objective matches a cost, like monetary cost to execute the query, query execution time, or energy consumption cost. Every alternative would then represent a single QEP.

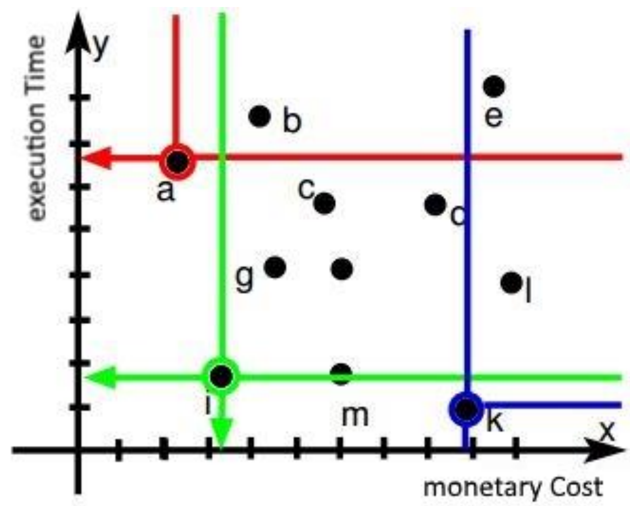

Figure 3: Pareto Optimal Solutions in a two-dimensional Query Optimization

Figure 3 gives an example for a two-dimensional query optimization. Each point represents a QEP where the points a, i, and k are not dominated by any other point in this graph.

The strength of finding a Pareto set is that every alternative in this set is optimal for at least one scoring function. A scoring function describes a specific stress configuration on the different objectives in order to set the importance to them and to compare alternatives in this Pareto set [9]. One disadvantage of a scoring function is that these stresses on objectives have to be defined prior to execution whereas the Pareto Set does not require any further input besides the QEPs. An advantage of a scoring function is that a user does not have to select a QEP out of a set since a specific scoring function only has one optimal solution. Figure 4 shows a scoring function: Every point in this two-dimensional space will be projected on a linear function using the stresses. The value of this projection is the score. We have proven, that the alternative with the minimum score is element of the Pareto Set [7].

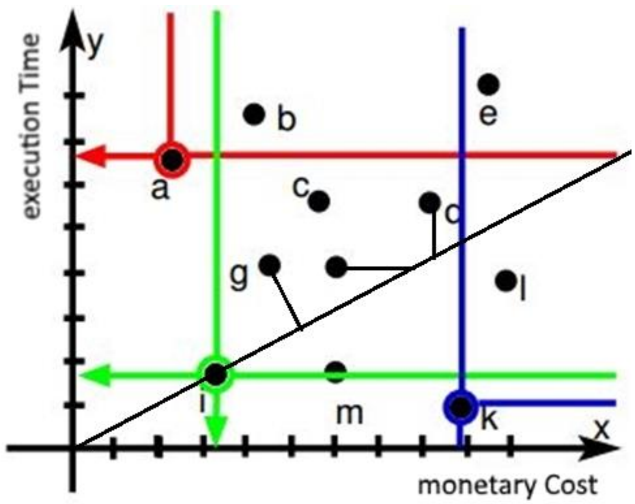

Figure 4: Scoring Function

In the field of Cloud Query Optimization, much research was done on Skyline Queries, which return the Pareto Set as their results [5]. Hence, the definition of a Skyline Query is similar to a Pareto Set: **"**The skyline query is to return a set of Pareto-optimal objects, called a skyline. Specifically, an object A can be a skyline object if there exists no other object B dominating A – an object A is said to dominate another object B if a scores better than B with respect to at least one utility function, and A does not score worse than B with respect to any other utility functions" [10]. Skyline queries were first mentioned in 2001 where three basic skyline computations were introduced [11]. Much research has been conducted to optimize the algorithm itself [12] [13]. Furthermore, important research was proposed by Trummer and Koch in 2014 which introduces parametric query optimization using a preprocessing of skyline queries to later on speed up the selection of a single solution via weights [14]. Unfortunately the preprocessing is limited to potentially relevant queries which first have to be defined, and the users are then restricted to those definitions.

However, the major drawback of a skyline query is the result of its: the skyline itself tends to be very large in size so that a user is left with his decision to select a single QEP out of a large pool of possible solutions [15]. This problem was also mentioned and addressed in the research to "Interactive Skyline Queries" in 2012 [10]: "However, the applicability of skyline queries suffers from a severe drawback because incomplete user preferences often lead to an impractical skyline size". This research introduces a way of reducing the skyline to a size manageable by a user. Users are stepwise asked about their preference between two chosen objectives. These questions are designed according to the answer's impact on the size of the skyline to minimize user interaction, but the approach still uses multiple user interactions.

Furthermore, the research on Skyline Queries neglects the fact that the user executing the query might not be aware of all constraints on the different costs to execute a query. In the use case of a larger organization, a query executing user might know the restrictions on execution time and energy consumption but might be unaware of the monetary cost he is allowed to spend since a budget is centrally managed by his supervisor. An approach is needed to separate query executing users from another type of users, whom we call superusers, who do not execute queries, but know and have authorities to set the restrictions on the queries.

In the following section, we briefly present our previous research on the Normalized Weighted Sum Algorithm (NWSA). Additionally, we also describe how this approach minimizes user interactions and enables a separation of query executing users from superusers.

3. NORMALIZED WEIGHTED SUM ALGORITHM

The Weighted Sum Model [16] is an existing optimization strategy which incorporates multiple objectives into its decision. Using a score for each alternative, which includes all objectives, an alternative is rated by a single number called score and can be compared to other alternatives. The score aggregates the different objectives, stressed by individual weights on each objective. Ordering the alternatives by the score allows the model to choose the best alternative: maximum score for utility functions and minimum score for cost functions. Used in many multi-objective optimization problems in various fields of computer science and others such as economics (Cost-Utility Analysis) [17] [18], the model lacks the idea of adding different dimensions and units leading to "adding apples and oranges" [19] if not addressed.

To use the Weighted Sum Model in the context of different dimension and unit objectives, the Normalized Weighted Sum Algorithm (NWSA) [7] uses the Weighted Sum Model as basis but makes major changes to cover the weaknesses.

$$A_i^{WSM-score} = \sum_{j=1}^{n} w_j \frac{a_{ij}}{m_j}$$

Figure 5: Modified Weighted Sum Model Scoring Function

Figure 5 shows the function used by NWSA to calculate the score of an alternative: normalizing objectives to a user-defined maximum eliminates units and results in a distribution on a percentage basis. The result of this normalization is then multiplied by a weight to individually stress objectives to the preference of a user. These strategies adapt the ideas of a user based decision [20].

$$w_j = \frac{uw_j * ew_j}{\sum(uw * ew)}.$$

Figure 6: Composite Normalized Weight Factor

Figure 6 shows a more sophisticated approach on using weights: using the given formula allows adjusting each user-defined weight by an environmental weight. An example for this usage would be a scenario where a high weight on energy consumption is selected at a time where the mobile device is fully charged and is connected to a power source. The system is then able to regulate the user weight on energy to focus more on other objectives.

Figure 7 shows the process for a user using Skyline Queries / Pareto Set (a) and NWSA (b), respectively. Comparing the both processes shows that NWSA requires the input on weights to stress the objectives but eliminates the step of a user decision based on the result of a skyline query.

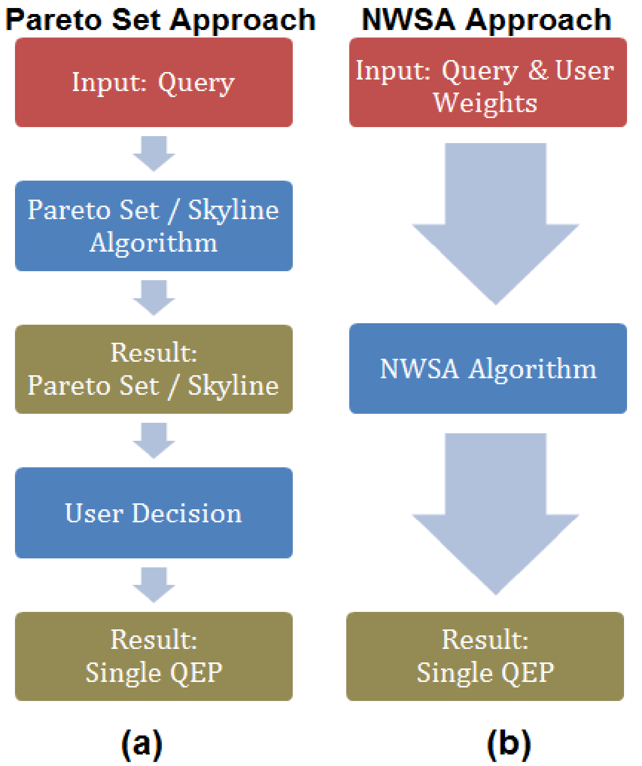

Figure 7: Query Optimization Processes: Pareto Set Approach (a) and NWSA Approach (b)

As discussed in the previous section, the final user decision on an alternative based on the result of the Skyline Query can still remain very complex. Research on the size of a Pareto Set from 1978 already estimated the size of its as $\Theta((\ln n)^{d-1} / (d-1)!)$ for n data objects and d objectives, assuming distinct value condition and attribute independence [6].

Analyzing the disadvantage NWSA has towards Skyline Queries that user preferences/weights have to be known prior to execution offers an advantage of being able to separate two types of users: users that make the decision on weights (superusers) and users that invoke the execution of a query (query executing users).

A set of weights, which we call a weight profile, can be preset by a superuser who is aware of all constraints on the different objectives. This weight profile can then be selected by a user executing the query, minimizing the decisions he has to make. Furthermore, a weight profile can be described by an application-based logical description, such as “emergency query” or “batch query”, to describe a weight profile with high importance on execution time or low importance on execution time, respectively. This abstraction to a simpler user interaction also reduces the complexity of the executing user’s decision.

Taking the advantage of being able to preset weight profiles by superusers and making the decision simple for executing users still leaves the superuser with his decision to set weights. A user interface can be given where previous data on decisions is used to estimate an impact on the overall costs of money, time, and energy when changing weights. Designing a sophisticated user interface for a superuser to predict the change of a weight is a major challenge since this interface has to display 6 dimensions: 3 weights and 3 costs. By adding a linear dependency between the weights, the dimensions can be broken down into three 2-dimentional graphs as the one shown in the example in Figure 8.

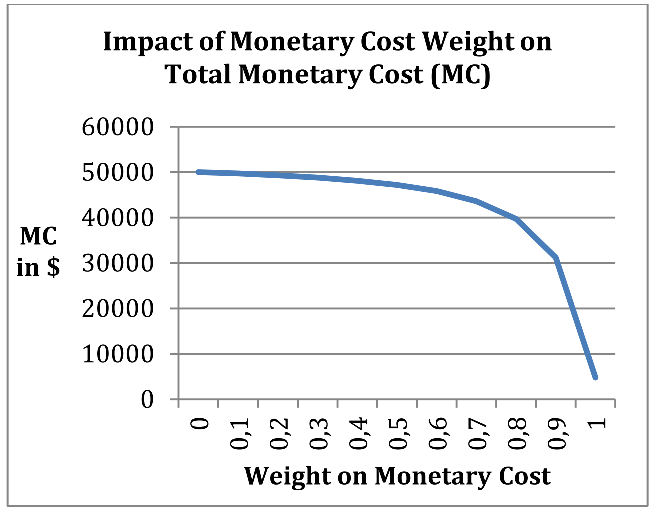

Figure 8: User-interface example for a superuser to adjust weights

4. EVALUATION

This section describes an evaluation of the explained difference between the Skyline query selection process and the weight choosing process needed by the NWSA. This comparison is based on a user study described in section 4.1 and followed by a discussion of the results in section 4.2.

An evaluation of the NWSA in terms of complexity, runtime and possible outcome of chosen QEPs in comparison to the one dimensional lexicographical ordering is described in the earlier publication of NWSA [7].

4.1 User Study

The goal of the user study is to compare the two approaches, the Skyline query selection and the weight selection process in NWSA, in terms of the accuracy of the decision that a user makes and the amount of time the user needs to make such a decision. The user study contains three sets of questions, each of which has four questions. The first set represents the decision a user has to make in the Skyline approach selecting an alternative based on the given Pareto Set. The second set represents the decision a user has to make in the NWSA approach selecting weights to stress the different objectives. The third set represents the idea of having a superuser preset a weight profile and describing the profile with a logical description. The sets appear in alternating order for different participants of the study to remove any bias towards any of the three approaches because of the order of the sets.

In each set, users are asked to select one alternative to buy a TV based on the given alternatives of each. The study has been transposed to this easy equivalent question so that no specific knowledge or large introduction to the field is needed. With each question in a set, users are presented with three, five, seven and nine alternatives to choose from. In Figures 9, 10, and 11, the examples of these different sets and questions are given.

User Behavior Studies for a Multi-Objective Problem Strategy

Evaluate the following options to buy a new TV under the following facts:

- Cost, delivery and trust are equally important to you as long as delivery is faster than 15 days and the cost is less than $750.

Select your option in the Drop-Down menu below.

Option #	Price	Shipping duration	Vendor Reputation 1(high) to 10(low)
Option 1	$500	10 days	4
Option 2	$600	9 days	2
Option 3	$350	14 days	6
Option 4	$800	5 days	7
Option 5	$300	18 days	10
Option 6	$350	15 days	5
Option 7	$700	20 days	1
Option 8	$550	10 days	3
Option 9	$900	2 days	8

N/A ▾

Proceed

Figure 9: User Study Set 1 representing the Skyline approach

User Behavior Studies for a Multi-Objective Problem Strategy

Evaluate the following options to buy a new TV under the following facts:

- You care more about a cheap option than a fast delivery
- You are risky and would rather trust an unknown vendor than wait on a longer delivery time.

Select your option in the Drop-Down menu below.

Option #	Price Preference	Shipping duration Preference	Vendor Reputation Preference
Option 1	3	7	2
Option 2	1	6	4
Option 3	5	4	3
Option 4	8	2	1
Option 5	6	1	5

N/A ▾

Proceed

Figure 10: User Study Set 2 representing the NWSA approach with weight profiles

User Behavior Studies for a Multi-Objective Problem Strategy

Evaluate the following options to buy a new TV under the following facts:

- The purchase cost is a little bit more important to you but it is not a big problem.

Select your option in the Drop-Down menu below.

Option #	Logical Description
Option 1	Express delivery for additional cost.
Option 2	Special deal from new vendor. Very fast delivery
Option 3	High reputation vendor with long waiting time.
Option 4	Buying from Craigslist: Very risky and might take a while to get a good offer.
Option 5	Overnight Express Delivery: A large additional cost but trusted vendor and very fast.
Option 6	Special deal from a trusted seller for an average price and an acceptable delivery time.
Option 7	Most expensive but fastest delivery and highest trust in seller.

N/A ▾

Proceed

Figure 11: User Study Set 3 representing the NWSA approach with logical descriptions

4.2 User Study Results

Preliminary results of the user study show that the logical description of a weight profile is by far the easiest decision out of the three given options. Participants in the study answered those decision questions in average nearly twice as fast (~42 seconds) as the decision with a given weight profile (~80 seconds). Participants needed similar time for selecting one solution out of the list of alternatives (~85 seconds). The accuracy of selecting the optimal answer was low for both the alternative list of study set one as well as for the weight profile of study set two (both < 50%). In contrast to that, participants selected the optimal alternative with accuracy greater than 80% given the logical descriptions from study set three.

Furthermore, giving a participant more than five alternatives to choose from in study set one and two increases the time needed to make a decision significantly (increase of ~40%) whereas participants needed an average of 10% more time to answer a question of seven or nine alternatives compared to five given alternatives in the case of study set three and the logical descriptions.

5. CONCLUSION AND OUTLOOK

This paper presented an interactive multi-objective query optimization approach based on the Normalized Weighted Sum Model (NWSM). The paper then presented a user study analyzing the user interaction with the query optimization process based on NWSA and based on) the general Skyline approach. In comparison to the Skyline approach, using NWSA eliminates the later decision process of the user on Pareto Optimal QEPs. This is achieved by selecting weights on the different objectives to automate the decision process on those Pareto optimal QEPs.

The results of our user study show, that using a logical description of a weight profile substantially increases the accuracy of a user selecting the optimal alternative and also speeds up the time a user needs to select his answer.

Our future research on NWSA includes enabling non-linear functions of the normalization of objectives as well as of the composition of user and environmental weights. As far as the usage of this algorithm is concerned, we intend to incorporate it into the query optimization process to calculate fast estimations of query costs for clouds [21, 22, 23, 24]. Another future area of usage of this algorithm is a new Cache Replacement Policy for the mobile Cache to extend semantic Caching [2]. Based on the computed score of a QEP, the new policy can help keep more valuable data in the semantic cache (the higher the score is, the higher the cost to regain the query data will be). Furthermore we are conducting research on a sophisticated interface for the superuser to set and adjust weight profiles according to his alternating budget requests.

ACKNOWLEDGEMENT

This work is partially supported by the National Science Foundation Award No. 1349285.

References

[1] J. Mullen, M. Perrin, F. Helff, L. Gruenwald and L. d'Orazio, "A Vision of Time-, Energy-, and Monetary Cost-Aware Query Processing in Mobile Cloud Database Environment," March 2015.

[2] M. Perrin, "Master Thesis at the University of Oklahoma: Time-, Energy-, and Monetary Cost-Aware Cache Design for a Mobile Cloud Database System," May 2015.

[3] Amazon, "Amazon EC2 Pricing," 2015. [Online]. Available: https://aws.amazon.com/ec2/pricing/?nc1=h_ls. [Accessed 18 08 2015].

[4] A. Ben-Tal, Characterization of Pareto and lexicographic optimal solutions, Springer Berlin Heidelberg, 1980.

[5] J. Chomicki, P. Ciaccia and N. Meneghetti, "Skyline queries, front and back," in *ACM SIGMOD Record 42.3*, 2013.

[6] J. L. Bentley, H. T. Kung, M. Schkolnick and C. D. Thompson, "On the average number of maxima in a set of vectors and applications," in *Journal of the ACM (JACM)*, 1978.

[7] F. Helff, L. Gruenwald and L. d'Orazio, "Weighted Sum Model for Multi-Objective Query Optimization for Mobile-Cloud Database Environments," in *Proceedings of the Workshops of the EDBT/ICDT 2016 Joint Conference EDBT/ICDT 2016*, Bordeaux, France, 2016.

[8] E. Zitzler and L. Thiele, "Multiobjective Evolutionary Algorithms:," in *IEEE TRANSACTIONS ON EVOLUTIONARY COMPUTATION, VOL. 3, NO. 4*, , NOVEMBER 1999.

[9] D. V. Lindley, Scoring Rules and the Inevitability of Probability, 50 ed., International Statistical Review Vol 50, 1982, pp. 1-11.

[10] J. Lee, G.-w. You, S.-w. Hwang, J. Selke and W.-T. Balke, "Interactive skyline queries," in *Information Sciences*, 2012.

[11] S. Borzsony, D. Kossmann and K. Stocker, "The skyline operator," in *17th International Conference on Data Engineering*, 2001.

[12] D. Kossmann, F. Ramsak and S. Rost, "Shooting stars in the sky: An online algorithm for skyline queries," in *Proceedings of the 28th international conference on Very Large Data Bases (VLDB)*, 2002.

[13] D. Papadias, Y. Tao, G. Fu and B. Seeger, "An optimal and progressive algorithm for skyline queries," in *Proceedings of the 2003 ACM SIGMOD international conference on Management of data (ACM)*, 2003.

[14] I. Trummer and C. Koch, "Multi-Objective Parametric Query Optimization," in *Proceedings of the VLDB Endowment*, 2014.

[15] Y. Tao, X. Xiao and J. Pei, "Efficient Skyline and Top-k Retrieval in Subspaces," in *TRANSACTIONS ON KNOWLEDGE AND DATA ENGINEERING*, 2007.

[16] E. Triantaphyllou, Multi-criteria decision making methods: a comparative study, 44 ed., Springer Science & Business Media, 2013.

[17] P. Kind, J. E. Lafata, K. Matuszewsk and D. Raisch, "The use of QALYs in clinical and patient decision-making: Issues and prospects," in *VALUE IN HEALTH 12*, 2012.

[18] M. G. M. Hunink and M. C. Weinstein, Decision Making in Health and Medicine: Integrating Evidence and Values, 2nd ed., Cambridge University Press, 2014.

[19] P. C. Fishburn, "Methods of estimating additive utilities.: Management science 13.7," 1967.

[20] C.-H. Goh, Y.-C. A. Tung and C.-H. Cheng, " A revised weighted sum decision model for robot selection," in *Computers & Industrial Engineering Vol.30(2)*, 1996.

[21] N. Bruno, S. Jain and J. Zhou, "Continuous cloud-scale query optimization and processing," in *Proceedings of the VLDB Endowment 6.11* , 2013.

[22] H. Kllapi, E. Sitaridi and M. M. Tsangaris, "Schedule Optimization for Data Processing Flows on the Cloud," in *Proceedings of the 2011 ACM SIGMOD International Conference on Management of data (ACM)*, 2011.

[23] M. Armbrust and e. al., "A view of cloud computing," in *Communications of the ACM 53.4*, 2010.

[24] H. M. Fard, R. Prodan, J. J. D. Barrionuevo and T. Fahringer, "A Multi-Objective Approach for Workflow Scheduling in Heterogeneous Environments," in *12th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing*, 2012.