CLAIMS

1. A database system comprising a computation platform, a search platform and a re-computation trigger platform, wherein

- the computation platform is arranged to
    - compute a plurality of database query results in response to each batch re-computation order received from either the search platform or the re-computation trigger platform, with the computation of the plurality of database query results for each batch re-computation order processed according to a batch processing schedule within a given time frame, wherein the timeframe is indicated in the batch re-computation order, determined autonomously by the computation platform, or negotiated between the computation platform and the search platform/re-computation trigger platform outside the re-computation orders;
    - return the computed database query results resulting from the batch re-computation orders to both the search platform and the re-computation trigger platform;
- the search platform is arranged to
    - initiate the batch re-computation orders based on a first update strategy;
    - maintain the computed database query results returned by the computation platform in a memory and
    - receive queries from clients connected to the search platform and in response to such queries, to make the computed database query results available to the clients; wherein
- the re-computation trigger platform comprises a trigger platform memory for holding data for determining when to initiate the re-computation orders to the computation platform, the data including a mirror of the computed database query results computed by the computation platform, and enriched with meta-data attributes, not part of the computed database query results returned by the computation platform, and wherein the re-computation trigger platform is further arranged to initiate the batch re-computation orders using a second update strategy which is different from the first update strategy.

2. The database system of claim 1 , wherein each batch re-computation order instructs the computation platform to compute at least one million database query results and within at most 24 hours.

3. The database system of claim 1 or claim 2, wherein the computation platform is further arranged to

- analyze and decompose each batch re-computation order into smaller computation packages relating to a subset of the plurality of database query results;
- assign the computation packages to processing modules in order to execute computations of the computation packages;
- re-assemble results of the computation packages computations performed by the processing modules, and
- return the re-assembled results to the search platform.

4. The database system of claim 3, wherein the arrangement of the computation platform to analyze and decompose each batch re-computation order includes an arrangement to

- split up each batch re-computation order into atomic computation transactions;
- identify redundant atomic computation transactions and discard them;
- group the remaining non-redundant atomic computation transactions to form the computation packages.

5. The database system of claim 3 or claim 4, wherein the computation platform is arranged to assign the computation packages to the processing modules according to the batch processing schedule, the batch processing schedule providing a load distribution by taking into account available computation resources of the computation modules within the given timeframe.

6. The database system of any one of the preceding claims, wherein, as part of the first update strategy, the search platform is arranged to

- assign to each of the computed database query results a score indicative of a needed refresh frequency;
- updating a portion of the computed database query results by directing a batch re-computation order to the computation platform, wherein the portion of the computed database query results to be updated is determined depending on the assigned score;
- in response to a query received by the search platform from a client, perform a database search to find those computed database query results which fulfil preferences included in the query;

- issue a response to the client and updating the scores assigned to the pre-computed database query results returned to the client.

7. The database system of any one of the preceding claims, wherein the search platform is further arranged to

- retrieve a plurality of computed database query results from the memory based on criteria in a search query received from one of the clients connected to the search platform;
- classify the retrieved computed database query results according to at least one category;
- rank the retrieved computed database query results within each category; and
- provide at least one of the ranked retrieved computed database query results for each category to display to the client.

8. The database system of any one of the preceding claims, wherein the re-computation trigger platform is either integrated into the computation platform or the search platform or being a separate entity, and as part of the second update strategy, the re-computation trigger platform is arranged to

- o determine probabilities of the mirrored computed database query results being outdated,
- o issue batch re-computation orders to the computation platform regarding a portion of the computed database query results having a higher probability of being outdated than other pre-computed database query results,

wherein the computation platform is further arranged to

- o receive batch re-computation orders from the re-computation trigger platform, and
- o return database computed query results resulting from the batch re-computation order to both, the search platform and the re-computation trigger platform.

9. The database system of claim 8 comprising a plurality of search platforms connected to the computation platform and the re-computation trigger platform is arranged to trigger re-computations of computed database query results populating anyone of the plurality of search platforms connected to the computation platform.

10. The database system of claim 8 or claim 9, wherein the re-computation trigger platform is further arranged to determine the probabilities of the mirrored computed database query results being outdated depending on a probabilistic model and on the occurrence of asynchronous real-time events, wherein the probabilistic model models discrepancies between the mirrored computed database query results maintained in the re-computation trigger platform and presumed actual database query results.

11. The database system of any one of claims 8 to 10, wherein the re-computation trigger platform is further arranged to

- analyse whether incoming asynchronous real-time events are represented in the probabilistic model;
- for real-time events which are not represented in the probabilistic model, issue re-computation orders regarding the respective particular mirrored pre- computed database query results as soon as possible;
- for real-time events which are represented in the probabilistic model, accumulate such real-time events over a certain period of time, compare the actually occurred and accumulated real-time events with their representation in the probabilistic model and, if the actually occurred accumulated real-time events deviate from their representation in the probabilistic model to a predetermined extent, issue a batch re-computation order with respect to potentially mirrored computed database query results as soon as possible.

12. The database system of any one of the preceding claims, wherein the database system is a travel search system or a travel reservation system, wherein

- the computation platform is arranged to have access to information on travel availability and fares and is arranged to compute priced travel recommendations on the basis of the travel availability information and fares;
- the search platform is arranged to maintain computed priced travel recommendations computed and received from the computation platform in response to batch re-computation orders and to return pre-computed priced travel recommendations to a client in response to a query, the returned pre- computed priced travel recommendations fulfilling preferences included in the query.

13. The database system of claim 12, wherein the search platform is arranged to

- retrieve a plurality of pre-computed priced travel recommendations from the memory based on criteria in a search query received from one of the clients connected to the search platform;
- classify the retrieved pre-computed priced travel recommendations according to at least one category;
- rank the retrieved pre-computed priced travel recommendations within each category; and
- provide at least one of the ranked retrieved computed pre-computed priced travel recommendations for each category to display to the client.

14. The database system of claim 13, wherein the category is at least one of an elapsed time for each of the pre-computed priced travel recommendations, a popularity of each of the pre-computed priced travel recommendations, a presence of an exclusive offer associated with each of the pre-computed priced travel recommendations, a cost of each of the pre-computed priced travel recommendations, a time at which each of the pre-computed priced travel recommendations was most-recently booked, or a sponsorship associated with each of the pre-computed priced travel recommendations.

15. A method for making pre-computed database query results available to clients, comprising:

- computing a plurality of database query results by a computation platform in response to each batch re-computation order received from either a search platform or the re-computation trigger platform, with the computation of the plurality of database query results for each batch re-computation order processed according to a batch processing schedule within a given time frame, and the computation platform returning the computed database query results resulting from the batch re-computation orders to the search platform, wherein the timeframe is indicated in the batch re-computation order, determined autonomously by the computation platform, or negotiated between the computation platform and the search platform/re-computation trigger platform outside the re-computation orders;
- initiating the batch re-computation orders by the search platform based on a first update strategy,

maintaining, by the search platform, the computed database query results returned by the computation platform in a memory,

- the search platform, receiving queries from clients connected to the search platform and in response to such queries, making the computed database query results available to the clients;
- the re-computation trigger platform, holding data for determining for determining when to initiate the re-computation orders to the computation platform in a trigger platform memory, the data including a mirror of the computed database query results computed by the computation platform, and enriched with meta-data attributes, not part of the computed database query results returned by the computation platform, wherein the method further comprises, the re-computation trigger platform, initiating the batch re-computation orders using a second update strategy which is different from the first update strategy.

16. The method of claim 15, wherein each batch re-computation order instructs the computation platform to compute at least one million database query results and within at most 24 hours.

17. The method of claim 15 or claim 16, further comprising the computation platform

- analyzing and decomposing each batch re-computation order into smaller computation packages relating to a subset of the plurality of database query results;
- assigning the computation packages to processing modules in order to execute computations of the computation packages;
- re-assembling results of the computation packages computations performed by the processing modules, and
- returning the re-assembled results to the search platform.

18. The method of claim 17, wherein the computation platform analyzing and decomposing the batch re-computation order further includes

- splitting up each batch re-computation order into atomic computation transactions;
- identifying redundant atomic computation transactions and discard them; and
- grouping the remaining non-redundant atomic computation transactions to form the computation packages.

19. The method of claim 17 or claim 18, wherein the computation platform further assigns the computation packages to the processing modules according to the batch processing schedule, the batch processing schedule providing a load distribution by taking into account available computation resources of the computation modules within the given timeframe.

20. The method of any one of claims 15 to 19, further comprising the search platform, as part of the first update strategy,

- assigning a score to each of the computed database query results, the score being indicative of a needed refresh frequency;
- updating a portion of the computed database query results by directing a batch re-computation order to the computation platform, wherein the portion of the computed database query results to be updated is determined depending on the assigned score;
- in response to a query received by the search platform from a client, performing a database search to find those computed database query results which fulfil preferences included in the query;
- issuing a response to the client and updating the scores assigned to the pre-computed database query results returned to the client.

21. The method of any one of claims 15 to 20, further comprising the search platform

- retrieving a plurality of computed database query results from the memory based on criteria in a search query received from one of the clients connected to the search platform;
- classifying the retrieved computed database query results according to at least one category;
- ranking the retrieved computed database query results within each category; and
- providing at least one of the ranked retrieved computed database query results for each category to display to the client.

22. The method of any one of claims 15 to 21, wherein the re-computation trigger platform is either integrated into the computation platform or the search platform or being a separate entity, and as part of the second update strategy, the method further comprises the re-computation trigger platform:

- determining probabilities of the mirrored computed database query results being outdated,
- issuing a batch re-computation order to the computation platform regarding a portion of the computed database query results having a higher probability of being outdated than other pre- computed database query results; and - the computation platform receiving the batch re-computation order from the re-computation trigger platform, and

- the computation platform returning computed database query results resulting from the batch re-computation order to both, the search platform and the re-computation trigger platform.

23. The method of claim 22, wherein a plurality of search platforms is connected to the computation platform and the re-computation trigger platform triggers re-computations of computed database query results populating anyone of the plurality of search platforms connected to the computation platform.

24. The method of claim 22 or claim 23, wherein the re-computation trigger platform further determines the probabilities of the mirrored computed database query results being outdated depending on a probabilistic model and on the occurrence of asynchronous real-time events, wherein the probabilistic model models discrepancies between the mirrored computed database query results maintained in the re-computation trigger platform and presumed actual database query results.

25. The method of any one of claims 22 to 24, further comprising the re-computation trigger platform

- analysing whether incoming asynchronous real-time events are represented in the probabilistic model;
- for real-time events which are not represented in the probabilistic model, issuing re-computation orders regarding the respective particular mirrored computed database query results as soon as possible;
- for real-time events which are represented in the probabilistic model, accumulating such real-time events over a certain period of time, comparing the actually occurred and accumulated real-time events with their representation in the probabilistic model and, if the actually occurred accumulated real-time events deviate from their representation in the probabilistic model to a predetermined extent, issuing a batch re-computation order with respect to potentially mirrored computed database query results as soon as possible.

26. The method of any one of claims 15 to 25, wherein the computation platform and the search platform form a travel reservation system, wherein

- the computation platform has access to information on travel availability and fares and computes priced travel recommendations on the basis of the travel availability information and fares;

- the search platform maintains pre-computed priced travel recommendations computed and received from the computation platform in response to batch re-computation orders and returns pre-computed priced travel recommendations to a client in response to a query, the returned pre-computed priced travel recommendations fulfilling preferences included in the query.

27. The method of claim 26, further comprising the search platform

- retrieving a plurality of pre-computed priced travel recommendations from the memory based on criteria in a search query received from one of the clients connected to the search platform;
- classifying the retrieved pre-computed priced travel recommendations according to at least one category;
- ranking the retrieved pre-computed priced travel recommendations within each category; and
- providing at least one of the ranked retrieved computed pre-computed priced travel recommendations for each category to display to the client.

28. The method of claim 27, wherein the category is at least one of an elapsed time for each of the pre-computed priced travel recommendations, a popularity of each of the pre-computed priced travel recommendations, a presence of an exclusive offer associated with each of the pre-computed priced travel recommendations, a cost of each of the pre-computed priced travel recommendations, a time at which each of the pre-computed priced travel recommendations was most-recently booked, or a sponsorship associated with each of the pre-computed priced travel recommendations.

29. A non-transitory computer readable storage medium having computer program instructions stored therein, which when executed on a computer system implementing the arrangements of the computation platform, the search platform and the re-computation trigger platform of any one of claims 1 to 14.