

METHOD AND SYSTEM FOR A PRE-SHOPPING RESERVATION SYSTEM WITH INCREASED SEARCH EFFICIENCY

Field of the invention

- 5 The present invention relates to the field of reservation systems, particularly to a method and system for a massive search platform increasing efficiency in travel searches.

Background of the invention

10

- State of the art reservation systems are normally based on dedicated Global Distribution Systems (GDS), as for example airlines reservation systems which provide flight search applications for shopping business like flight booking. This activity, also called “shopping for booking” involves a lot of computation and may
- 15 take some time. To minimize this delay, users usually have few degrees of freedom: they must specify origin and destination cities, outbound and inbound dates of the journey. To further minimize the delay, they may specify e.g. a preferred operating carrier and cabin class if the have precise flights requirements. Users search for the best price for a particular travel (city pair,
- 20 departure and arrival date) in the aim of booking it eventually. The search usually offers some flexibility: e.g. returning the 100 cheapest flights recommendations for the requested travel; returning cheaper flights for closely-related dates. All needed computation (searching the cheapest fares and rules combination, checking seat availability of candidate flights...) are performed at the time of the
- 25 query, which ensures that the returned recommendations will be available for booking. Consequently, such search transactions are costly and take several seconds to complete. This cost precludes them from answering more open search requests, such as for instances the cheapest flight for the coming two or three months. While this is advantageous for system performance and for
- 30 response times, it is not ideal for users who would certainly appreciate a more user friendly interaction with wider freedom in the parameters choice.

A different approach to the task of searching air travel prices is the so called “pre-shopping”. With this term we refer to those activities which require interrogations of data bases through a reservation system but which do not necessarily result in a proper booking. This activities are of key importance for airlines or travel agencies, because, even if they do not generate an immediate revenue they can influence the future choice of their potential customers. It would be highly appreciated a tool able to provide a zero-delay response to a user’s query with many degrees of freedom. With Pre-shopping, users can browse a carrier or travel agency’s entire catalogue of air travels. Those users wish to make their mind prior to shopping by browsing recommendations over billions of travels recommendations. Compared to shopping, browsing recommendations implies instantaneous responses to searches (few tens of milliseconds). The typical approach of pre-shopping systems is thus to let users browse a cache of pre-computed travel recommendations. With such approach, the search queries can be much more powerful: users can search on many open criteria: origin city only, range of dates, range of price... For the sake of the example:

“Search for 2 or 3 weeks trips from Paris to any destination within the next 12 month below 600 Euros”

The drawback of this approach is that the recommendations returned to users are only guaranteed to be valid at the time of their pre-computation. In particular, they may no longer be eligible for booking at the time of the search.

Unlike other cache browsing domains (e.g. WWW search), air travel pre-shopping is very sensible to air travel price volatility: the best prices of flights in the coming weeks is likely to change every day. This volatility greatly impacts the cache accuracy, i.e., the consistency between price in pre-shopping and price in shopping. The usual accuracy rate in the industry is about 20-30%.

Maintaining higher cache accuracy often means massive re-computations (to deal with the entire catalogue of travel) and also frequent re-computations (to deal with flight volatility). This is very demanding in hardware resources.

3

State of the art pre-shopping tools have some drawbacks which limit the efficiency of the tool. For example TravelTainment pre-shopping platform: ("TTibe: TravelTainment Internet Booking Engine" <http://www.traveltainment.fr/a-propos-de-traveltainment/qui-sommes-nous/>) provides browsing facilities over its

5 own database of pre-computed travels (flights departing from German cities mainly). The air travel data are provided e.g. by Amadeus' Extreme Pricer, a product of Massive Computation Platform (MCP). Travel data represent the cheapest flights from several thousands of city pairs, for every day of the coming year, for all stay durations between 1 and 23 days. Every day, the entire base of
10 travel (several tens of millions of prices) is recomputed by Amadeus and sent to TravelTainment for integration into their platform. While the travel domain is rather exhaustive from its customers' standpoint, this approach has two main drawbacks:

-All data are recomputed by Amadeus every day, which has an operational
15 cost

- Integrating this amount of data is costly for TravelTainment and can only be performed once a day. This has an impact on price accuracy experienced by their customers.

20 Other commercially available platforms are Kayak's Explore (<http://www.kayak.com/news/kayak-adds-map-based-search-tool-to-popular-ipad-app.bd.html>),

Opodo's EscapeMap:

(http://promos.opodo.co.uk/airtools/escape_map.html),

25 Lufthansa's Trip Finder:

(<http://www.lufthansa.com/online/portal/lh/us/nonav/local?nodeid=3322431&l=en>)

which is powered by Amadeus technologies. These three pre-shopping platforms have a different strategy than TravelTainment to feed their cache of pre-computed solution: they all rely on recording real shopping traffic, i.e. record the
30 result of search transaction operated on their shopping platform. This approach

has an advantage in that the pre-computation comes at almost no cost. However it comes with a series of penalties for their respective customers:

- unusual destinations might not be available to pre-shopping due to recorded traffic;

- 5 - there are a lot of "holes" in the price domains, due to missing dates in recorded traffic

- it is difficult to propose complex travel recommendations, for example advance purchase

- 10 - some volatile price recommendations are not updated for days (even weeks).

All these disadvantages can compromise to a great extent the pre-shopping accuracy experienced by the customers.

15 **Object of the invention**

An object of the present invention is to alleviate at least some of the problems associated with the prior art systems.

20 According to one aspect of the present invention there is provided a method in a distributed reservation system for generating priced travel recommendations according to non-binding travel queries, the distributed reservation system having access to a plurality of travel databases containing information on travel availability and fares according to a plurality of parameters, each travel query

25 including a set of preferences each preference being related to a parameter selected among the plurality of parameters, the method including: maintaining on a plurality of fast access memory location a cache including a selection of pre-computed travel recommendations each travel recommendation including information on fares and/or availability for a specific travel, sorted by at least one

30 of the plurality of parameters; assigning to each of the cached travel recommendations a score indicative of a needed refresh frequency; updating the

selection of pre-computed travel recommendations by launching a massive query in the plurality of databases for refreshing the information included in at least some of the travel recommendations, wherein the frequency of the refresh of each travel recommendation depends on the assigned score; responsive to a travel query being received by the system, searching the plurality of fast access memory locations, to find those travel recommendations which fulfil the preferences included in the travel query; issuing a response to users for the travel query and updating the score of the selection of travel recommendations.

10 According to a second aspect of the present invention there is provided a system comprising one or more components adapted to perform the method described above.

15 According to a further embodiment of the present invention there is provided a computer program comprising instructions for carrying out the method described above when said computer program is executed on a computer system.

The method according to a preferred embodiment of the present invention allows storing entire catalogues of air travels from many providers, while ensuring a high cache accuracy at limited operational costs. The system makes use of the fares knowledge to merge identical travel recommendations coming from different travel providers when possible (i.e. when fares are public and not negotiated for a specific customer). This prevents redundant prices to be stored in the system and improves its cost effectiveness. The system according to a preferred embodiment of the present invention relies on several dedicated data analysis engines to optimize the cost of data pre-computations while at the same time maintaining good data accuracy.

Brief description of drawings

30 Reference will now be made, by way of example, to the accompanying drawings, in which:

Figure 1 is a diagram of the massive search platform for a pre-shopping reservation system in accordance with one embodiment of the present invention;

- 5 Figure 2 shows a possible reservation system implementing the massive search platform of the present invention;

- Figures 3-8 show the details of the single module components of the massive search platform for a pre-shopping reservation system in accordance with one
10 embodiment of the present invention;

Figure 9 is a diagram of a general computer system adapted to support the method of a preferred embodiment of the present invention;

- 15 Figure 10 is a flow chart of the method steps of a process, in accordance with one embodiment of the present invention.

Detailed description of the embodiments

- 20 The distributed search platform according to a preferred embodiment of the present invention aims at storing pre-computed air travel prices and provides several advantages.

- It is designed to hold entire catalogs of air travel. The search platform is able to store best prices of e.g. thousands of markets for all the days of the year to
25 come, several possible prices per day depending of the travel's stay duration. The distributed nature of the invention lets it scale to whatever number of markets to store.

It optimizes storage of air travel data, with the following benefits:

- Travel search efficiency: the achievable response time of search
30 transactions is *few tens of milliseconds*, whatever their complexity (open destination, full year search...)

- Travel data update efficiency: update has a limited impact on the overall storage, which enables continuous updates, with instantaneous availability of new data.
- Travel data storage efficiency: data can be factorized if common in many catalog to further reduce storage cost

It can maintain a high quality of pre-shopping travel recommendations at a sustainable cost. The system has various engines to detect the travels prices that require re-computations. This can drive partial re-computation to save hardware resource. The saved resource can be re-invested in other re-computations to achieve higher cache accuracy from a user's standpoint (in the range of 80-90%).

The system provides different types of search products, depending on the needs of its customers and travel providers. For the sake of the example, a carrier would need a search product to retrieve recommendations for his airline and that of its partners. On the contrary, an online travel agency would need a search product to retrieve any type of air travel without carrier filtering. Those two products may have internal specificities and may be optimized accordingly.

Figure 1 shows a distributed system 101 for storing and searching air travel recommendations according to a preferred embodiment of the present invention. The stored catalogues or air travel recommendations are spread over all the machines which compose the distributed system. On every machine, the subpart of the global data is stored in a fast access memory location (e.g. RAM) for fast data access. Figure 1 presents a logical view of the architecture of a system implementing the present invention. Although physically distributed, the logical view of the system can be split in two main groups:

The data bases 103 and square boxes 105 and 107 represent the typical parts of a distributed search system.

Databases 103 represent all the air travel recommendations logically grouped into pools of recommendations and physically stored across different machines

8

The square boxes 105 and 107 represent the Feed and Search engines:

- The Feed engine 105 indexes groups of pre-computed air travel recommendations (e.g. all travels with same city pair) and stores them in one of the machine in the distributed system.
- 5 - The Search Engine 107 locates groups of data among the physical machines and provides index search facilities to answer to queries originating from users.

The oval items 109-115 represent a series of business-oriented analysis engines.

- 10 Their purpose is to optimize the hardware cost of the platform (and thus the cost for the travel providers): they aim at achieving a good compromise between the number of recommendations to re-compute every day vs. the accuracy of the pre-computed prices stored in the system. These engines analyze feed and search operations and generate metrics on volatility and quality of data stored in
- 15 the system. Some of those engines make use of other shopping services of the GDS (not part of the invention). In particulars:

- a Learning Engine 109, analyses the travel recommendations which are stored in the platform every day to extract some information about the volatility of some prices across dates and markets, i.e. how long a
- 20 price stays the same;
- a Popularity Engine 111 tracks the most requested or most returned travel recommendations, (per date, market...) to gain metrics concerning the most relevant data stored in the system.
- an Accuracy Engine 113 tries to detect discrepancies between cached
- 25 travel recommendations stored in the system and real shopping prices, i.e. flights which are no longer bookable or those whose prices have changed
- a Report Coordinator 115 aggregates and stores the results of the
- 30 previous engines. Its role is to determine which part of the entire flights data must be recomputed given available resources and based on the

results of the previous engines. Such determination is done according to an algorithm.

In the system, all the analysis engines work in parallel to Feed and Search Engines (105 and 107), so their work has no performance impact for the users
5 (no response time degradation).

With reference to Figure 2 a complete Search Platform is represented. Such platform includes a system 201 according to a preferred embodiment of the present invention as one of the component which is fed by a travel
10 recommendations computing system 203, such as the one described in the European Patent Application EP11305518.0 (in the following description we will refer to this system as Massive Computation Platform, or MCP). Both systems may interact with other internal services provided by a GDS 205.
Figure 2 depicts the high level flow to follow in order to store entire catalogs of
15 travel recommendation in the search platform.

As customers of the platform, the travel providers (airline, travel agency...) decide which part of their travel domain they want to integrate into the search platform. From that point, they send to the travel recommendations computing system a
20 so-called massive query that is a series of computation orders to the travel recommendations computing system. Those order details the markets to consider (e.g. a list of city pairs for all days of the year to come) as well as the travel recommendations to generate (e.g. for every day, the best recommendations for journeys between 1 and 20 days long). Such orders can be
25 re-evaluated frequently by the customer or they can serve as a base for a recurring computation.

The travel recommendations computation system makes use of internal services of the GDS to compute the requested recommendations. Among other things, to
30 generate a recommendation, it may use journey services to retrieve the list of

existing flights; pricing services to find the best combination of fares and flights; availability services to consult the current seats available for booking...

As the recommendations are generated, the computing system sends the results
5 to the system according to a preferred embodiment of the present invention for integration. The received recommendations are stored in dedicated memory locations to populate a global cache of pre-computed travel recommendations, becoming available for eventual users' search queries. Once travel recommendations are integrated, some monitoring tasks take place in
10 background to detect cached travel recommendations which must be recomputed due to a low consistency with equivalent shopping recommendations. This monitoring may use internal services provided by the GDS.

When inconsistent recommendations are detected, the system according to a
15 preferred embodiment of the present invention generates a series of computation orders (massive queries) and sends them to the computing system. The latter will generate recent recommendations which will help in maintaining a good cache consistency.

20 Figure 3 represents the structure and the functions of the Feed Engine 105 (see Figure 1). The Feed engine receives groups of air travel recommendations returned by a travel recommendations computing system, e.g. MCP. For example, all prices from Paris to London (market PAR-LON). The integration of those data occurs in several steps:

25

- Dispatch travel recommendations

In the system according to a preferred embodiment of the present invention, related data aims at being stored on the same physical machine to deliver very fast search response time. For example, two markets with the same origin city
30 (PAR-LON and PAR-NYC) will land on the same physical machine.

The Feed engine extracts information from the group of recommendations (travel provider ID, office ID, market, geographic location...) to determine the physical machine(s) which will host the data. This data balancing mechanism makes use of well known distributing techniques such as round robin, or consistent hashing.

- 5 As seen in well known data replication schemes, many machines can host the same group of recommendations to improve reliability, fault-tolerance, or accessibility.

- Organize travel recommendations

- 10 The Feed Engine receives batch of travel recommendations from a travel recommendations computing system, e.g. MCP. The incoming data are then grouped into what is called data sets, in a manner that suits better the system according to a preferred embodiment of the present invention. Each search product provided by the presently described system has a unique data
- 15 organization strategy to optimize its performance. For the sake of the example, for a particular need a group of flights coming from the travel recommendations computing system, e.g. MCP, could be organized in groups of identical city pairs and then be assigned to two types of data sets: 1) same city pair and direct flights only; and 2) same city pair and flights with connections.

20

- Build accelerators

- On top of that organization, the system according to a preferred embodiment of the present invention creates additional data sets that contain only meta-information about travel recommendations. These data help achieving very fast
- 25 searches. For example, a data set of meta-information can host the cities reachable from an origin city and for each reachable city the cheapest price for city pair. Eventually, searches could benefit from this information to avoid looking at too many travel recommendations before returning solutions.

- 30 - Build indexes

Like databases, the system according to a preferred embodiment of the present invention constructs indexes on top of data sets to provide fast access time. For pre-computed air travel recommendations, the searches criteria are very open: one can search prices in a given range of date, for a given range of price, for
5 arbitrary destination cities... Instead of creating one index per search criteria, the system according to a preferred embodiment of the present invention uses a well known multi-dimensional data structure (the K-D-B-tree) to construct a single index per data set, while maintaining an equally efficient access to data whatever the search criteria. This approach limits the amount of storage used.

10 If two travel providers share common travel recommendations, and the fares are public (contrary to negotiated fares of travel providers, only applicable to specific office ID), those can be shared in the system storage to gain space and reduce hardware cost.

15 - Consistency manager

When new travel recommendations are available from the travel recommendations computing system, e.g. MCP, the equivalent data sets are updated with new or less travel recommendations, and their index is rebuilt.

20 Concurrently, the impacted data sets may be searched for travel recommendations.

To prevent impacting on-going searches (both in term of performance and consistency), the Feed Engine manages revisions of data sets. While searches are being performed on revision n of a data set, a feed constructs revision $n+1$.

25 When the revised data set is constructed and indexed it becomes the new reference data set for searches. The previous data set is deleted from the storage memories of all the physical machines in the distributed system which hosts the data set, shortly after. This ensures those data are kept available for on-going searches until they finish and thus prevent consistency issue

Figure 4 represents the search engine 107 (see Figure 1). The search engine receives air travel search requests from users and crawls into all the data to return the best recommendation. This process occurs in several steps:

5 - Server affinity

The incoming search request must be processed by a specific search product provided by the system according to a preferred embodiment of the present invention. From that point, it must be routed to a physical machine that contains the data necessary for answering the request. The air travel recommendations
10 were dispatched by the Feed Engines to physical machines based on the specificities of the search product the recommendations were aimed at. The Search Engine performs the opposite operation: it analyses the search query to answer to, and based on its type it determines the data to search and the physical machine(s) where they are located.

15

 - Determine search domain

Once the search request is forwarded to the relevant physical machine, the search engine determines the data set where to find meta-information related to the search query. The meta-information is used to locate all the data sets of air
20 travel recommendations which contain potential solutions to the search query.

 - Search execution

Once all the potential air travel recommendation are located, the search engine parses all the multi-dimensional indexes to collect the best travel
25 recommendations based on the criteria expressed in the search query, e.g., the price range, the date range, the flight carrier etc...

25

The search engine can take advantage of the previous meta-information to decide not to search into all potential travel solutions. For the sake of the example, suppose the search query asked for the best price departing from a
30 specific city NCE. Suppose during the search a travel for destination city PAR was found for 100euro. If the meta-information states that the lowest prices for

city NYC is 500euro, the search engine will not even try to search solutions from NCE to NYC. This further decreases the response time of the search transaction.

- Related Searches

5 In case the Search Execution step returned no travel recommendations, the user's query may be too restrictive: the reasons for the lack of match are thus analyzed for each city pair previously considered. As an example, reasons may be that no flight exists in the specified date range, or flights exist but are more expensive than the limit expressed in the query.

10 In case the user's query is too restrictive, the Search Engine implements a fallback strategy to return recommendations which relates closely to the constraints expressed in the original query. It relaxes the constraints on the plurality of parameters (wider date range, higher price limit...) and loops back to a search execution step. The loop ends either when some recommendations are
15 found, when a configured number of retry is reached or when a maximum allowed time for retry is reached.

In case the fallback strategy does not return any recommendation, another fallback strategy is implemented when applicable. In case the requested destination has a geographic meaning (e.g., city, country), the Search Engine
20 uses the geographic services provided by the GDS (not part of the invention) to determine close geographic regions, it widens the destination constraints of the original query and loops back to the search execution step, in the manner explained above.

If both fallback strategies fail to retrieve recommendations, an empty result is
25 returned.

- Travel solutions aggregation

Once all solutions to the search query are found, the search engine performs a pass to merge identical results which could have been returned from different
30 data sets. This case can arise for example if the search had to look into a data

set containing the cheapest direct flights for a city pair, and in another data set containing all flights (direct and with stop).

At last, the found travel solutions are organized based requirements of the search query: group solutions by destination city, by date, by ascending price, by airline (either direct or with connections). The result is then returned to the user.

Figure 5 shows the Learning Engine 109 (See Figure 1). Learning Engine includes a Learning and statistics Database. The purpose of the Learning Engine is to detect the air travels whose prices do not change frequently and thus which do not need frequent re-computation to stay accurate.

The Learning Engine bases its logical analysis on a general property of air travel flights: the price of a flight scheduled in the coming days is very volatile, i.e., if the same flight (same departure date) is priced again a day after, its price has likely changed. On the contrary, the price of a flight scheduled several months away is unlikely to change if it is price again a day after or a week after.

The Learning Engine associates a volatility model to each market based on the property detailed above. It maintains the volatility model (and rectifies it if needed) based on the travel recommendation it receives every day.

- Record price of travel recommendations

When incoming travel recommendations are received, they are duplicated and one copy goes to the Learning Engine. Travel recommendations are grouped by market, i.e., recommendations for travels for one city pair, for days of the coming year. Note that not all days of the year yield recommendations, because the system might instruct the travel recommendations computing system (e.g. MCP) to recompute only volatile flights across small range of dates.

The Learning engine extracts the prices of the incoming travel recommendations and record them in a dedicated Learning Database, along with their date of

computation. Those prices serve as a basis for price comparison for future air travel data integrations.

- Load market data and rectify volatility model

- 5 The Learning Engine loads all the prices previously stored in his dedicated database for the incoming market. It compares the saved prices with the incoming prices available.

The Learning Engine adapts the volatility model for the market depending of the outcome of prices comparison:

- 10 When two identical flights have different prices, the difference is stored as a statistics in the Learning Database. When differences occur too frequently, the volatility model is updated: the span date range is marked more volatile. Storing statistics about change frequency helps mitigate the effect of price changes due to punctual events, like holidays season.
- 15 If two identical flights have the same price for a longer period than what is expected based on the model, the model is also updated: the span date range is marked as less volatile.

- Generate volatility reports

- 20 Once the analysis of all the incoming travel recommendations is finished, the Learning Engine generates a report which contains the revised volatility of the data which have just been integrated in the distributed search platform described herein. A volatility report is generated per customer ID (provider of the data) and is organized per market, per departure date.

25

The generated reports are sent to another Engine called the Report Coordinator. The latter will eventually use this source of information to decide of the subset of the air travel recommendations which must be recomputed depending on the available computing resource on the travel recommendations computing system

- 30 (e.g. MCP).

Figure 6 shows the Popularity Engine 111 (See Figure 1). Popularity Engine includes a Popularity Database. On incoming search requests, the Popularity Engine gets a chance to extract search trends out of users transactions. This gives insights about the popularity of travel prices stored in the system.

5 The Analysis is performed in steps:

- Input and output analysis

Before reaching the search engine, the input search query is duplicated and one copy goes into the Popularity Engine. Symmetrically, the output of the search
10 transaction is duplicated before being sent back to the user and a copy goes to the Popularity Engine.

The Popularity Engine analyses both input and output of a transaction to gain some popularity metrics. The analysis yields different information depending on the criteria of the input search query. For example:

15 If the search query requested travel recommendations for a specific market (city pair), the engine can extract from the query a ranking on popular departure dates per market.

If the search query requested travel recommendations from a single origin city, the engine can extract from the solution a ranking on preferred destination cities
20 (or destination countries) per originating city and per price, date range...

- Storing statistics in database

The trends extracted from the input query and output solutions are stored on a dedicated Popularity Database. This storage is by nature distributed, so that any
25 physical machine (where the Popularity Engine operates) can benefit from the data produced on other physical machines of the system.

- Aggregate Records and generate Popularity Reports

In parallel, a recurring job of the Popularity Engine is to extract the statistics
30 previously computed from the distributed database and to build some systemic popularity metrics.

For example, based on the total number of search queries that were analyzed for popularity, extract a ranking of popular markets, i.e., markets which were most targeted by input queries, cheaper markets returned in output travel solutions...

Another possibility is to generate statistics about most popular market for given
5 date ranges throughout the year, to extract trends for specific events like holiday season or generate statistics about most popular markets for different geographic zones in the world. This refinement helps extract more relevant popularity measures (e.g., on domestic flights...)

The generated reports are then sent to the Report Coordinator to give him
10 access to popularity metrics.

Figure 7 shows the Accuracy Engine 113 (See Figure 1). Accuracy Engine includes a Price Accuracy Database.

The goal of the Accuracy Engine is to control pre-computed air travel
15 recommendations and detect those whose price diverges too much from the real (i.e., current) shopping price. The principle of the engine is to use external Amadeus shopping services (not part of the invention) to shoot shopping transactions and to compare the returned prices with that of the cache of travel recommendations.

20

This engine operates in several steps.

-Use input and output transactions, with throttling

Like for the Popularity Engine, the Accuracy Engine receives a copy of the
25 duplicated input search query and output travel solutions.

To avoid shooting too many real shopping transactions, which are expensive in hardware cost and in response time, a throttling pass is applied to operate on a subset of the natural traffic that goes through the distributed search platform described herein.

30

- Generate Fare Search transactions

One must ensure that the Accuracy Engine will generate fare search transactions which are diversified enough to analyze a representative subset of the travel recommendations stored in the system. To do so, the generation strategy is twofold:

5 Generate fare search transactions based on travel popularity: the Accuracy Engine accesses the Popularity Database (presented earlier) to analyze the popularity of output solutions and keep only the most popular ones for further analysis. This maximizes the consistency experienced by users regarding cached prices vs. real shopping prices.

10 Generate random transactions based on the contents of the cache of travel recommendations. A random election of travels for analysis aims at providing eventual accuracy of the entire cache of travel recommendations. It ensures good travel findability for the users, i.e., accuracy of unpopular flights is also monitored, though less often to limit hardware cost.

15

- Aggregate data

The gathered accuracy metrics are stored in a dedicated, distributed accuracy database. A recurring job consolidates all the accuracy metrics into several reports (accuracy of solutions by customer, by market, by departure date...) and
20 sends them to the Report Coordinator for further usage.

Figure 8 shows the Report Coordinator 115 (See Figure 1). The Report Coordinator is the last business analysis engine of the system according to a preferred embodiment of the present invention. Its role is to decide which part of
25 the data in the platform must be re-computed, based on all the metrics reported by the previous engines and according to the available computing resource on the travel recommendations computing system (e.g. MCP).

The decision of re-computation is performed in several steps.

30

Store Reports

All the metrics from the incoming volatility, popularity and accuracy reports are stored locally in a dedicated report database. This storage is necessary because the Report Coordinator takes its decision based on previous metrics. For example, the Report Coordinator infers the age of data stored in the cache of travel recommendations based on the volatility reports. This information is then kept in the Report Database.

Re-computation decision

The decisions made by Report Coordinator are based on heuristics to balance data accuracy of the cache vs. computation resource on the computing system (e.g. MCP).

The basic approach is to re-compute the least accurate cached travel recommendations, given the available resource on the computing system (e.g. MCP). Among the least accurate data, the most popular ones are considered first for the generation of a re-computation order.

The candidates for re-computation are selected by the Report Coordinator in order to form groups of nearby travels in the flights domains (close date range for each market). This allows the travel recommendations computing system to mutualise some re-computations and further optimize its resource consumption.

Between each re-computation order, the Report Coordinator makes use of the volatility models (stored in the report database) and the inferred age of the travel recommendations to update the forecast accuracy of all the remaining data in the MSP.

With reference to Figure 9 a generic computer of the system (e.g. computer, reservation server, data base management subsystem, router, network server) is denoted with 950. The computer 950 is formed by several units that are connected in parallel to a system bus 953. In detail, one or more microprocessors control operation of the computer 950; a RAM 959 is directly used as a working memory by the microprocessors 956, and a ROM 962 stores basic code for a

bootstrap of the computer 950. Peripheral units are clustered around a local bus 965 (by means of respective interfaces). Particularly, a mass memory consists of a hard-disk 968 and a drive 971 for reading CD-ROMs 974. Moreover, the computer 950 includes input devices 977 (for example, a keyboard and a mouse), and output
5 devices 980 (for example, a monitor and a printer). A Network Interface Card 983 is used to connect the computer 950 to the network. A bridge unit 986 interfaces the system bus 953 with the local bus 965. Each microprocessor 956 and the bridge unit 986 can operate as master agents requesting an access to the system bus 953 for transmitting information. An arbiter 989 manages the granting of the
10 access with mutual exclusion to the system bus 953. Similar considerations apply if the system has a different topology, or it is based on other networks.

Alternatively, the computers have a different structure, include equivalent units, or consist of other data processing entities (such as PDAs, mobile phones, and the like).

15 The method described above is also represented in the diagram shown in figure 10. The method is a pre-shopping tool (i.e. for generating priced travel recommendations according to non-binding travel queries) which operates in a distributed reservation system having access to a plurality of travel databases
20 containing information on travel availability and fares according to a plurality of parameters, each travel query including a set of preferences each preference being related to a parameter selected among the plurality of parameters. The distributed reservation system includes a plurality of fast access memory locations where a selection of travel recommendations is maintained so that
25 users can have a fast response to their queries. We refer to these fast access memory with the term of "cache made of pre-computed travel recommendations" even if those skilled in the art will appreciate that this term is technically not completely appropriate because, as opposed to a cache, the contents stored in the fast access memories represent entire travel catalogues and not only a
30 subset comprising of the most accessed travels of the catalogues. In other words, the term cache relates to "shopping transactions cache". Similarly, the

information stored on those fast access memories are referred to as "cache information" (e.g. "cached travel recommendations"). The problem with the cache of pre-computed travel recommendations is that the data contained in the recommendations need to be updated in order for their price to keep consistent with the real prices found during shopping. This is a very costly (in term of time and system performances) activity and a balance must be found between consistency of information and system performances. One of the main features of the present invention is that the different travel recommendations are not updated at the same time and with the same frequency; rather a "frequency" score is assigned to each travel recommendation, according to which the update is scheduled: the information which is likely to be more volatile will be refreshed more frequently. The method begins at black circle 1001 and then goes to box 1003 where the system maintains a cache made of a selection of pre-computed travel recommendations on a plurality of fast access memory locations, each travel recommendation including information on fares and/or availability for a specific travel, sorted by at least one of the plurality of parameters. The sorting and indexing of the cached travel recommendations can be done in several ways as explained above: this is important to identify the right information in the shortest time possible. A score indicative of a needed refresh frequency is then assigned to each of the cached travel recommendations (step 1005). This score will be used to determine when a refresh of the information should be made. The refresh (step 1009) according to a preferred embodiment of the present invention is done with a batch processing as explained above: in a preferred embodiment of the present invention a massive query is launched through a dedicated subsystem, e.g. a dedicated travel recommendations computing system as for example the Massive Computation Platform mentioned above. The launching of such massive query (step 1007) can be done at fixed time or invoked by a user or being triggered by specific events: for example the massive query could be done every pre-determined period of time (e.g. at the end of day or every hour); it could be automatically performed when a critical mass of queries is received or when the maximum capacity is reached; or again it could be requested by the

administrator of the system or by the travel providers. When a user query is received (step 1011) the system, tries first the search within the fast access memory locations (1013). If no solution is found (step 1015) then an optional fall-back action could be launched (step 1017): if no fall-back action is provided a message to the user will advise that no results are available in the system. Alternatively a new search could be launched with adjusted parameters, considering that the user's query may be too restrictive: the reasons for the lack of match are thus analyzed for each city pair previously considered. As an example, reasons may be that no flight exists in the specified date range, or flights exists but are more expensive than the limit expressed in the query. In case the user's query is too restrictive, the Search Engine implements a fallback strategy to return recommendations which relates closely to the constraints expressed in the original query. It relaxes the constraints on the plurality of parameters (wider date range, higher price limit...) and loops back to a search execution step. The loop ends either when some recommendations are found, when a configured number of retry is reached or when a maximum allowed time for retry is reached. In case the fallback strategy does not return any recommendation, another fallback strategy is implemented when applicable. In case the requested destination has a geographic meaning (e.g., city, country), the Search Engine uses the geographic services provided by the GDS (not part of the invention) to determine close geographic regions, it widens the destination constraints in of the original query and loops back to the search execution step, in the manner explained above. If both fallback strategies fail to retrieve recommendations, an empty result is returned. Yet another possible solution would be to launch a query to external databases to find the requested travel recommendation and to enrich the cache of pre-computed travel recommendations. If a result is obtained the cache of pre-computed travel recommendations will be enriched with this new information. In any case a response is issued to the user (step 1019). The score of the travel recommendations might need an update as well. In the embodiment described herein the queries are sent by users looking for pre-shopping information, i.e.

information on e.g. trip availability, fares, time or general information not necessarily aimed at completing a reservation. In a preferred embodiment of the present invention the system receiving the queries and performing the database enquiries for satisfying user queries is separate from the actual reservation
5 system, but those skilled in the art will appreciate that the two systems (pre-shopping and reservation) could be integrated together.

It will be appreciated that alterations and modifications may be made to the above without departing from the scope of the disclosure. Naturally, in order to
10 satisfy local and specific requirements, a person skilled in the art may apply to the solution described above many modifications and alterations. Particularly, although the present disclosure has been described with a certain degree of particularity with reference to preferred embodiment(s) thereof, it should be understood that various omissions, substitutions and changes in the form and
15 details as well as other embodiments are possible; moreover, it is expressly intended that specific elements and/or method steps described in connection with any disclosed embodiment of the disclosure may be incorporated in any other embodiment as a general matter of design choice.

Similar considerations apply if the program (which may be used to
20 implement each embodiment of the disclosure) is structured in a different way, or if additional modules or functions are provided; likewise, the memory structures may be of other types, or may be replaced with equivalent entities (not necessarily consisting of physical storage media). Moreover, the proposed solution lends itself to be implemented with an equivalent method (having similar
25 or additional steps, even in a different order). In any case, the program may take any form suitable to be used by or in connection with any data processing system, such as external or resident software, firmware, or microcode (either in object code or in source code). Moreover, the program may be provided on any computer-usable medium; the medium can be any element suitable to contain,
30 store, communicate, propagate, or transfer the program. Examples of such medium are fixed disks (where the program can be pre-loaded), removable disks,

tapes, cards, wires, fibres, wireless connections, networks, broadcast waves, and the like; for example, the medium may be of the electronic, magnetic, optical, electromagnetic, infrared, or semiconductor type.

In any case, the solution according to the present disclosure lends itself to
5 be carried out with a hardware structure (for example, integrated in a chip of semiconductor material), or with a combination of software and hardware.

CLAIMS

1. A method in a distributed reservation system for generating priced travel
5 recommendations according to non-binding travel queries, the distributed
reservation system having access to a plurality of travel databases containing
information on travel availability and fares according to a plurality of parameters,
each travel query including a set of preferences each preference being related to
a parameter selected among the plurality of parameters, the method including:

10 - maintaining on a plurality of fast access memory location a cache
including a selection of pre-computed travel recommendations each travel
recommendation including information on fares and/or availability for a specific
travel, sorted by at least one of the plurality of parameters;

- assigning to each of the cached travel recommendations a score
15 indicative of a needed refresh frequency;

- updating the selection of pre-computed travel recommendations by
launching a massive query in the plurality of databases for refreshing the
information included in at least some of the travel recommendations, wherein the
frequency of the refresh of each travel recommendation depends on the assigned
20 score;

- responsive to a travel query being received by the system, searching the
plurality of fast access memory locations, to find those travel recommendations
which fulfil the preferences included in the travel query;

- issuing a response to users for the travel query and updating the score of
25 the selection of travel recommendations.

2. The method of claim 1 wherein the score indicative of the needed refresh
frequency includes a value indicative of the volatility of the information on
availability and/or fares of the travel recommendation.

3. The method of any preceding claim wherein the score indicative of the needed refresh frequency includes a value indicative of the popularity of the travel recommendation.

5 4. The method of any preceding claims wherein the access to the plurality of travel databases and the step of enriching the selection of pre-computed travel recommendations is done by means of a dedicated sub-system.

5. The method of any preceding claim further comprising:

10 - dividing the selection of pre-computed travel recommendations into a plurality of homogeneous groups, wherein each travel recommendation in an homogeneous group has at least a predetermined number of parameters with the same value, and storing each group in one of the fast access memory locations.

15 6. The method of any preceding claim wherein the step of launching the massive query is done at predetermined time interval.

7. The method of any preceding claim wherein the massive query is done with a batch process.

20

8. The method of any preceding claim wherein the travel recommendations stored on the plurality of fast access memory locations are indexed according to a multi-dimensional data structure.

25 9. The method of any preceding claim further including:

- responsive to no results being obtained by searching the plurality of fast access memory locations, launching a series of closely related queries with relaxed constraints on the plurality of parameters, in the aims of returning approaching results for information purposes.

30

10. The method of any claim 1 to 8, further including:

- responsive to no results being obtained by searching the plurality of fast access memory locations, issuing an alert message to the user.

11. A computer program comprising instructions for carrying out the steps of
5 the method according to any preceding claim, when said computer program is executed on a computer.

12. A computer program product including computer readable means
10 embodying the computer program of claim 10.

13. A software tool for pre-shopping system including the method of any claim
1 to 10.

14. A distributed data processing system for managing pre-shopping travel
15 queries, the distributed data-processing system having access to a plurality of travel databases containing information on travel availability and fares according to a plurality of parameters, each travel query including a set of preferences each preference being related to a parameter selected among the plurality of parameters, wherein the system comprises one or more components adapted to
20 perform the method of any claim 1 to 10.

15. A service deployed in a data processing system for implementing the method of any claim 1 to 10.