



(19) **United States**

(12) **Patent Application Publication**
Tracy et al.

(10) **Pub. No.: US 2014/0257881 A1**

(43) **Pub. Date: Sep. 11, 2014**

(54) **DEMAND FORECASTING SYSTEMS AND METHODS UTILIZING FARE ADJUSTMENT**

Publication Classification

(71) Applicant: **US AIRWAYS, INC.**, Tempe, AZ (US)

(51) **Int. Cl.**
G06Q 10/02 (2006.01)

(72) Inventors: **Jonathan Tracy**, Tempe, AZ (US); **James Fox**, Tempe, AZ (US); **Julianne Bowles**, Tempe, AZ (US); **Michael Costello**, Tempe, AZ (US); **Daniel Muzich**, Phoenix, AZ (US); **Thomas Trenga**, Mesa, AZ (US)

(52) **U.S. Cl.**
CPC *G06Q 10/02* (2013.01)
USPC **705/5**

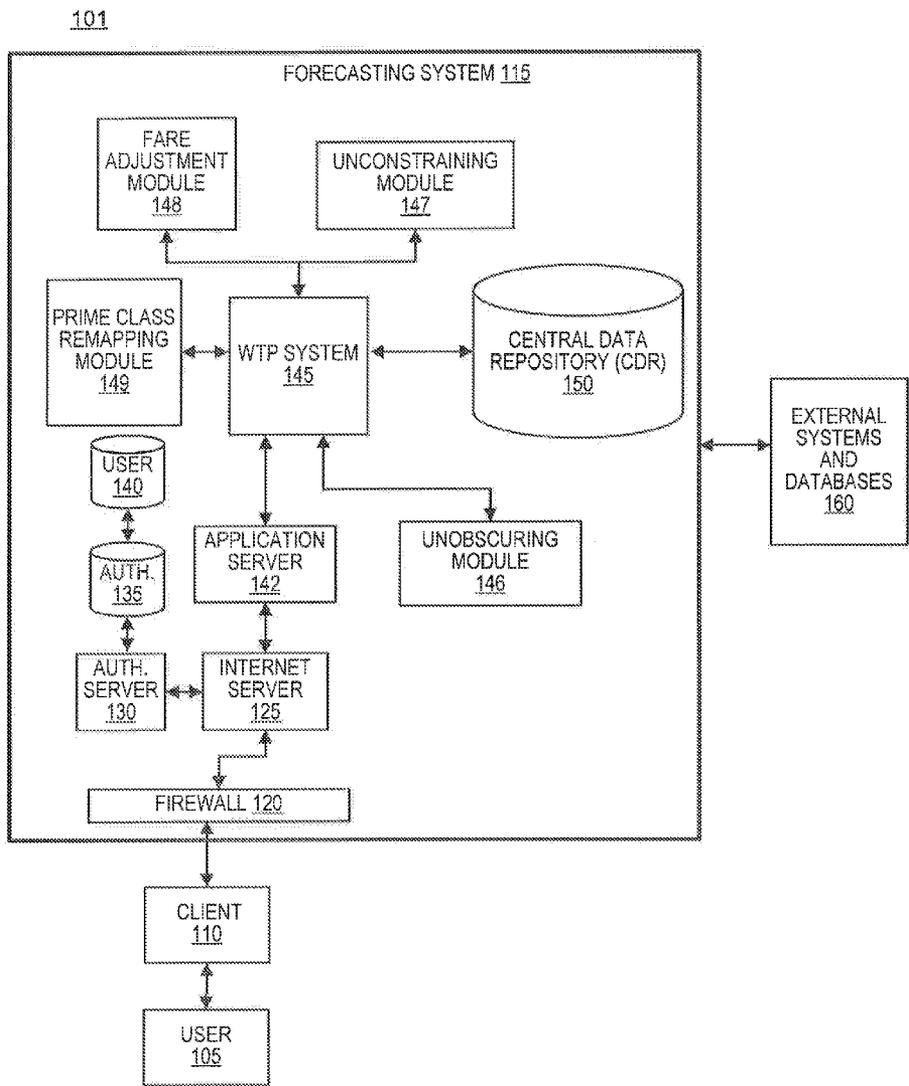
(73) Assignee: **US AIRWAYS, INC.**, Tempe, AZ (US)

(57) **ABSTRACT**

A computer based system for fare adjustment is disclosed. For a particular airline fare class, the system evaluates the financial impact of opening or closing the class, for example taking into account both stimulation effects and dilution effects. Via use of an exemplary fare adjustment system, decision-makers may increase revenues (and/or decrease losses) associated with airline fare class openings and closings.

(21) Appl. No.: **13/791,691**

(22) Filed: **Mar. 8, 2013**



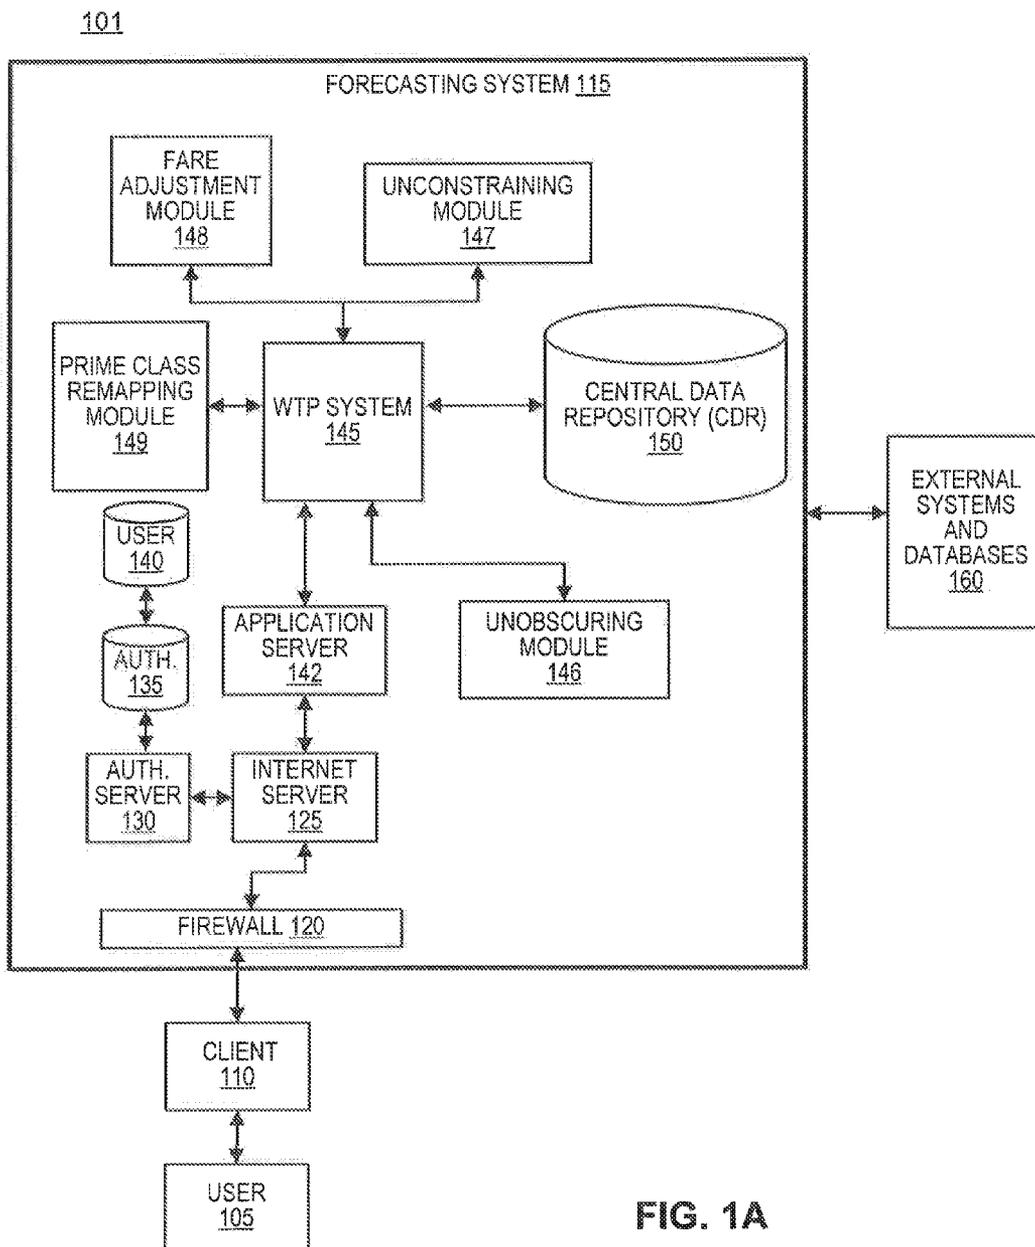


FIG. 1A

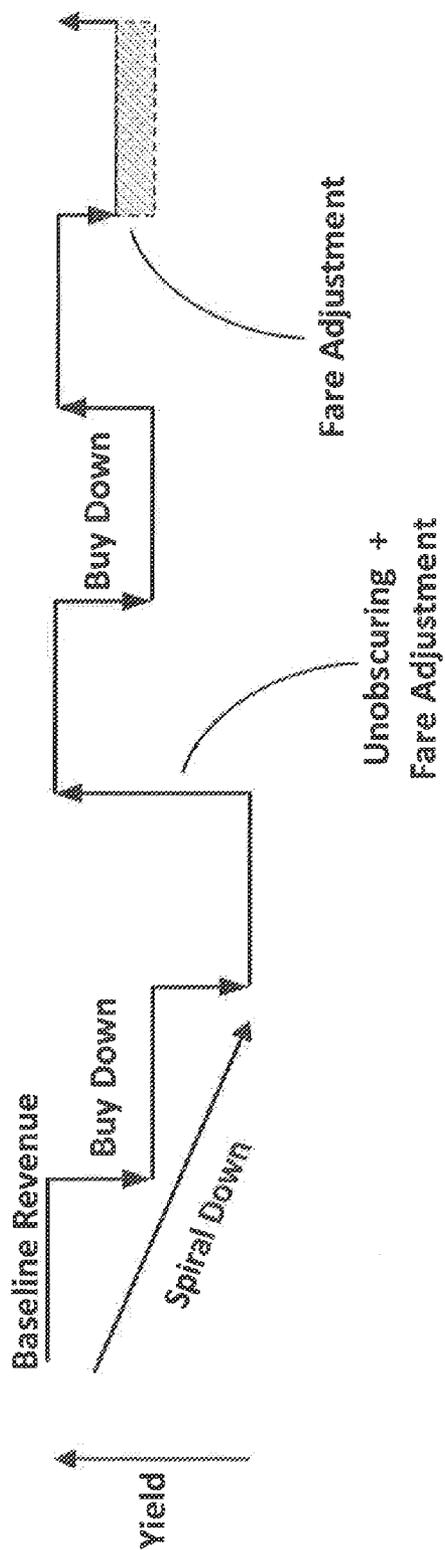


FIG. 1B

200

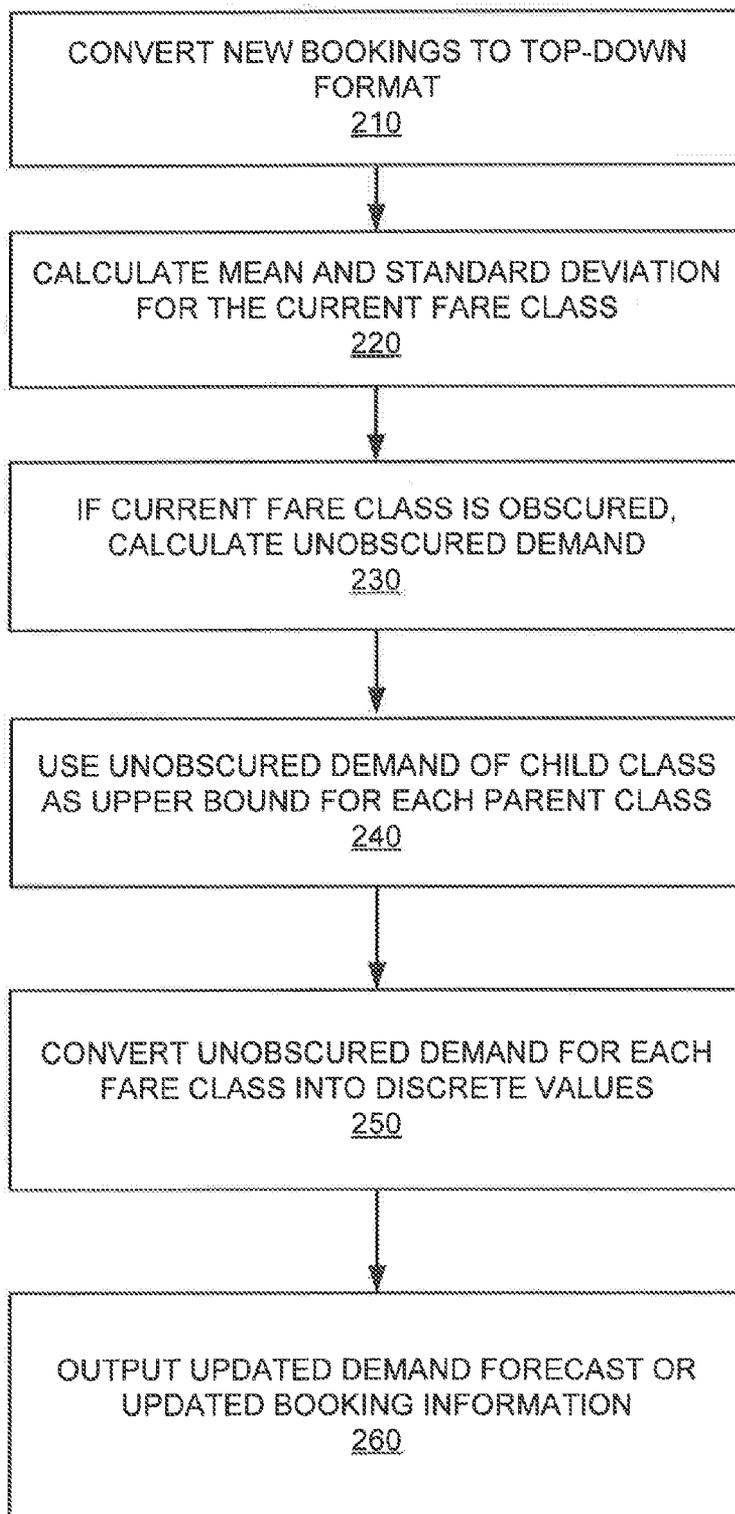


FIG. 2A

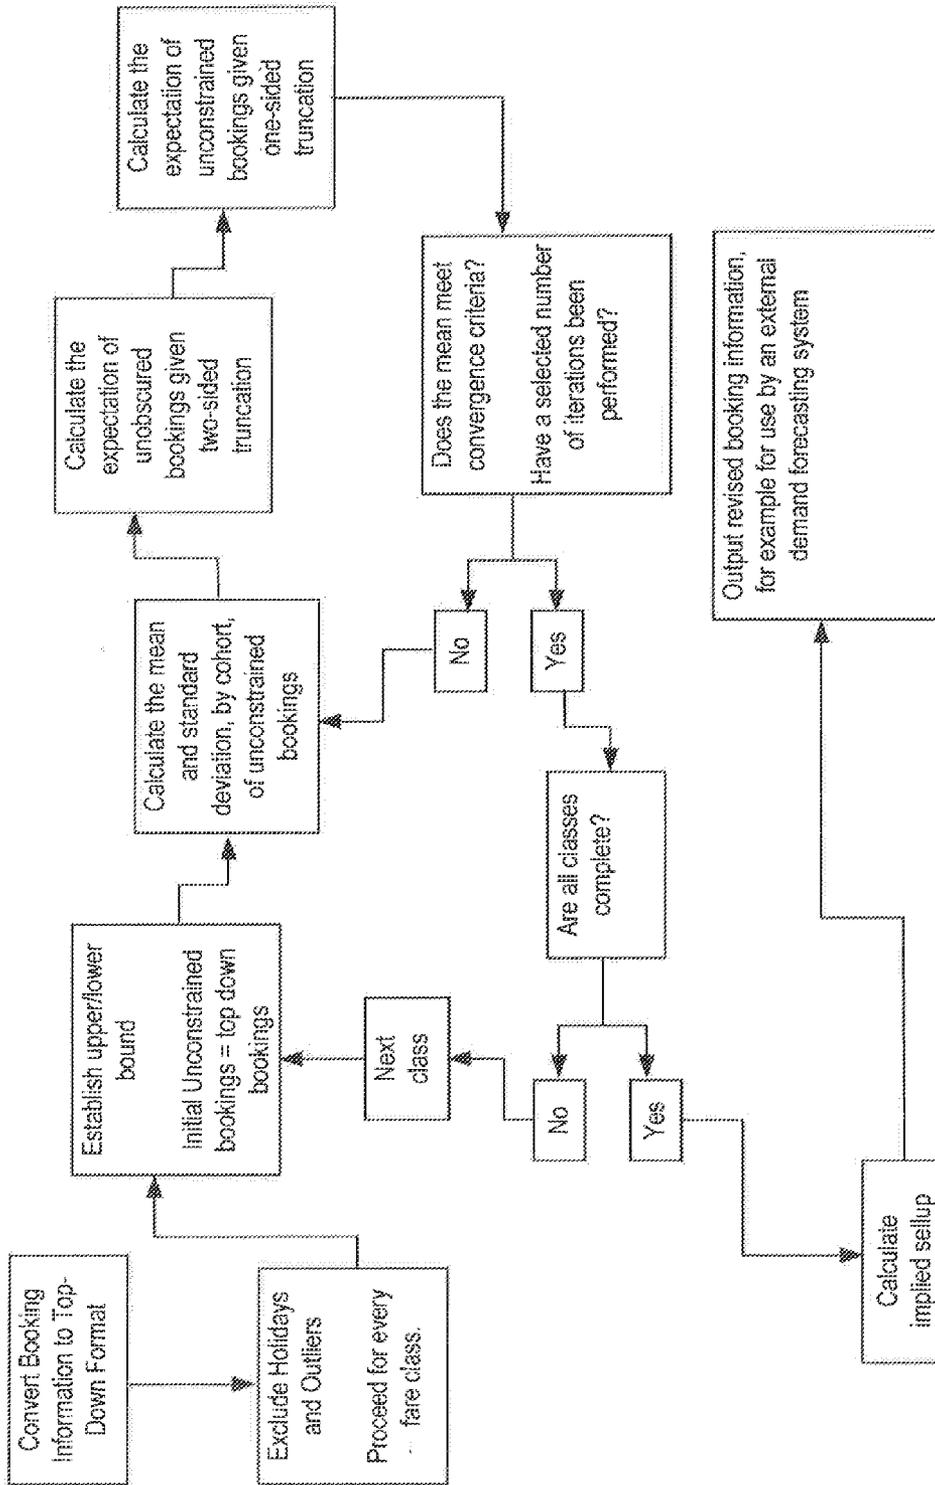


FIG. 2B

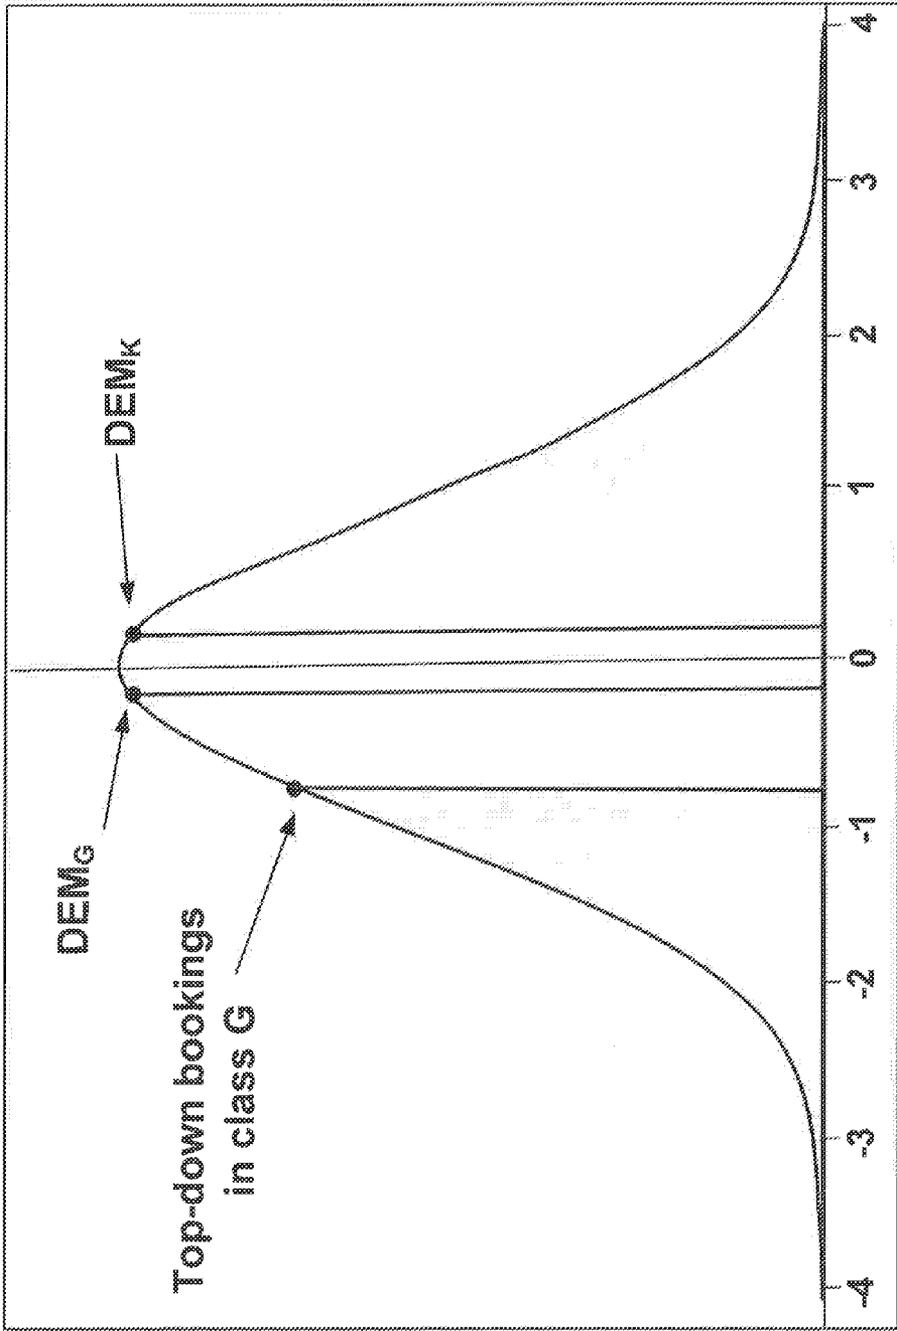


FIG. 2C

300

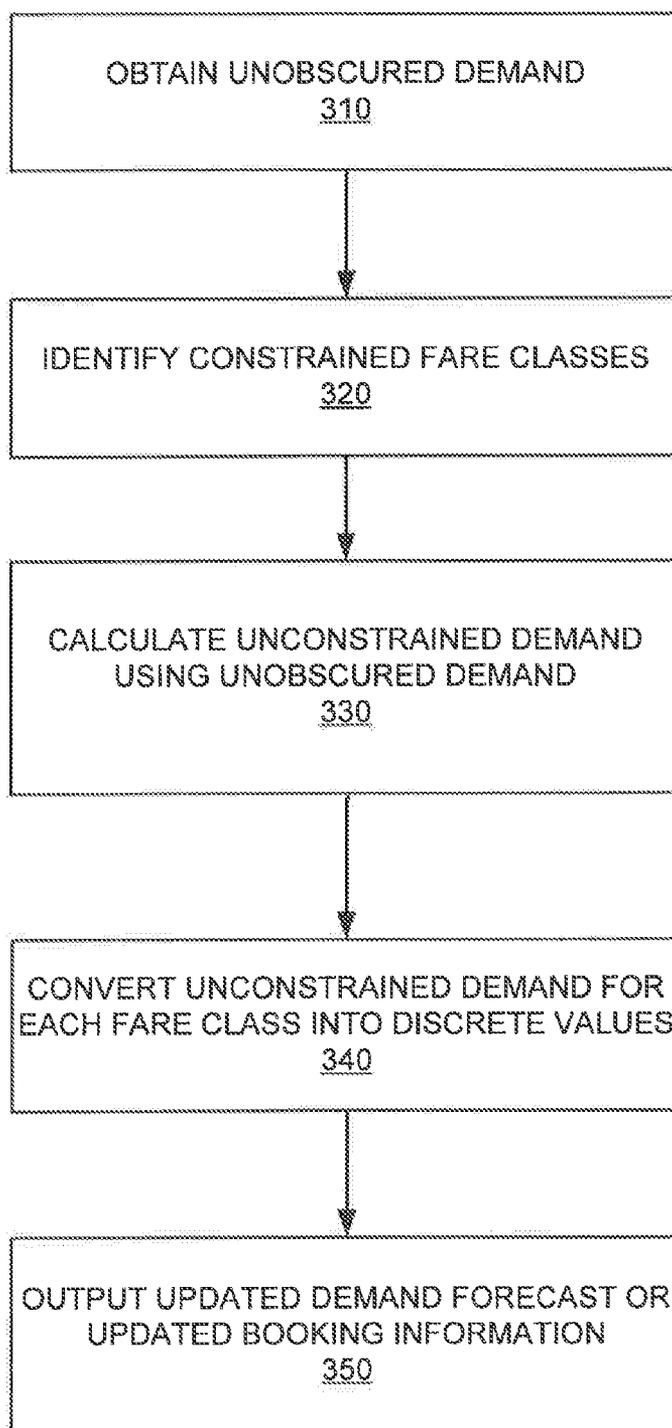


FIG. 3A

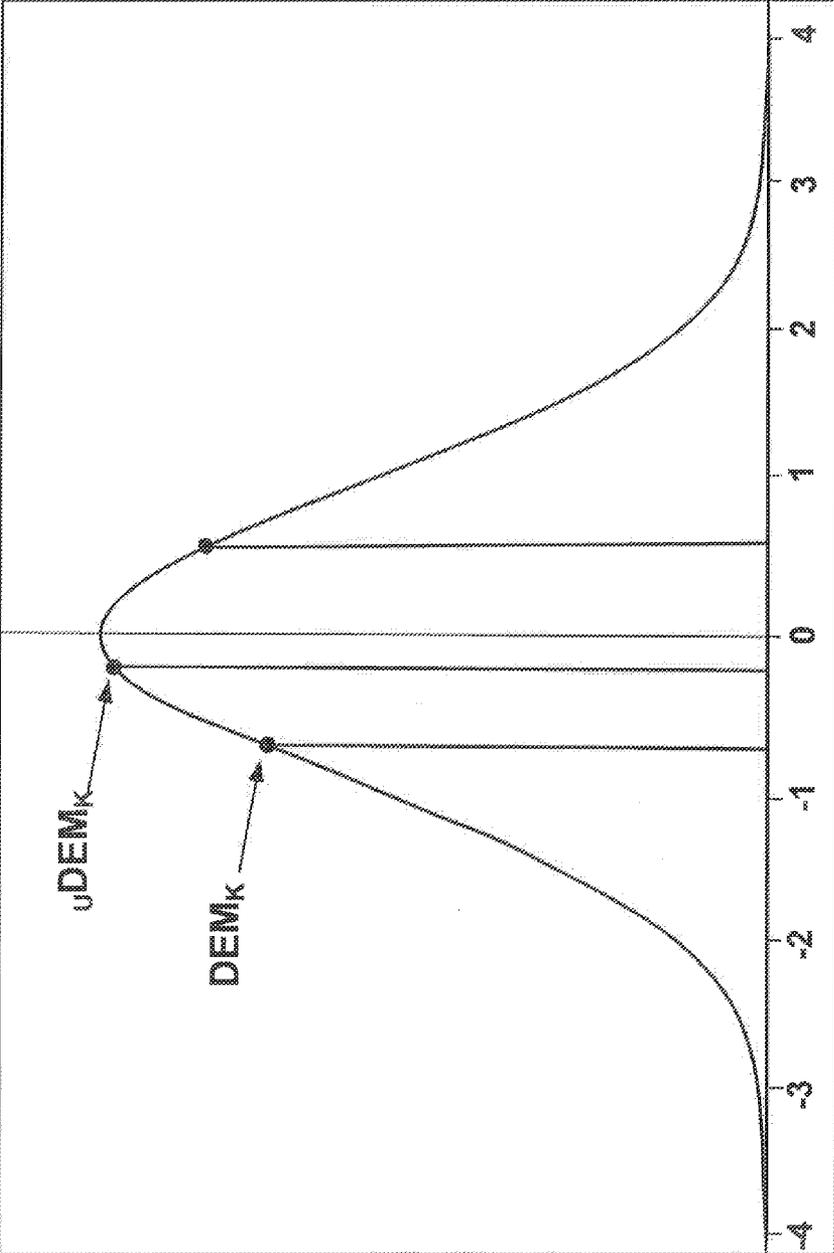


FIG. 3B

400

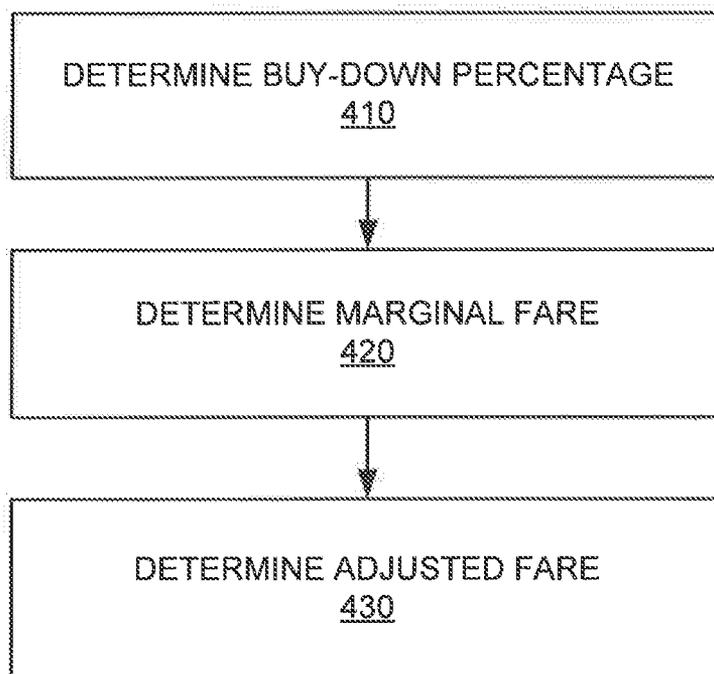


FIG. 4A

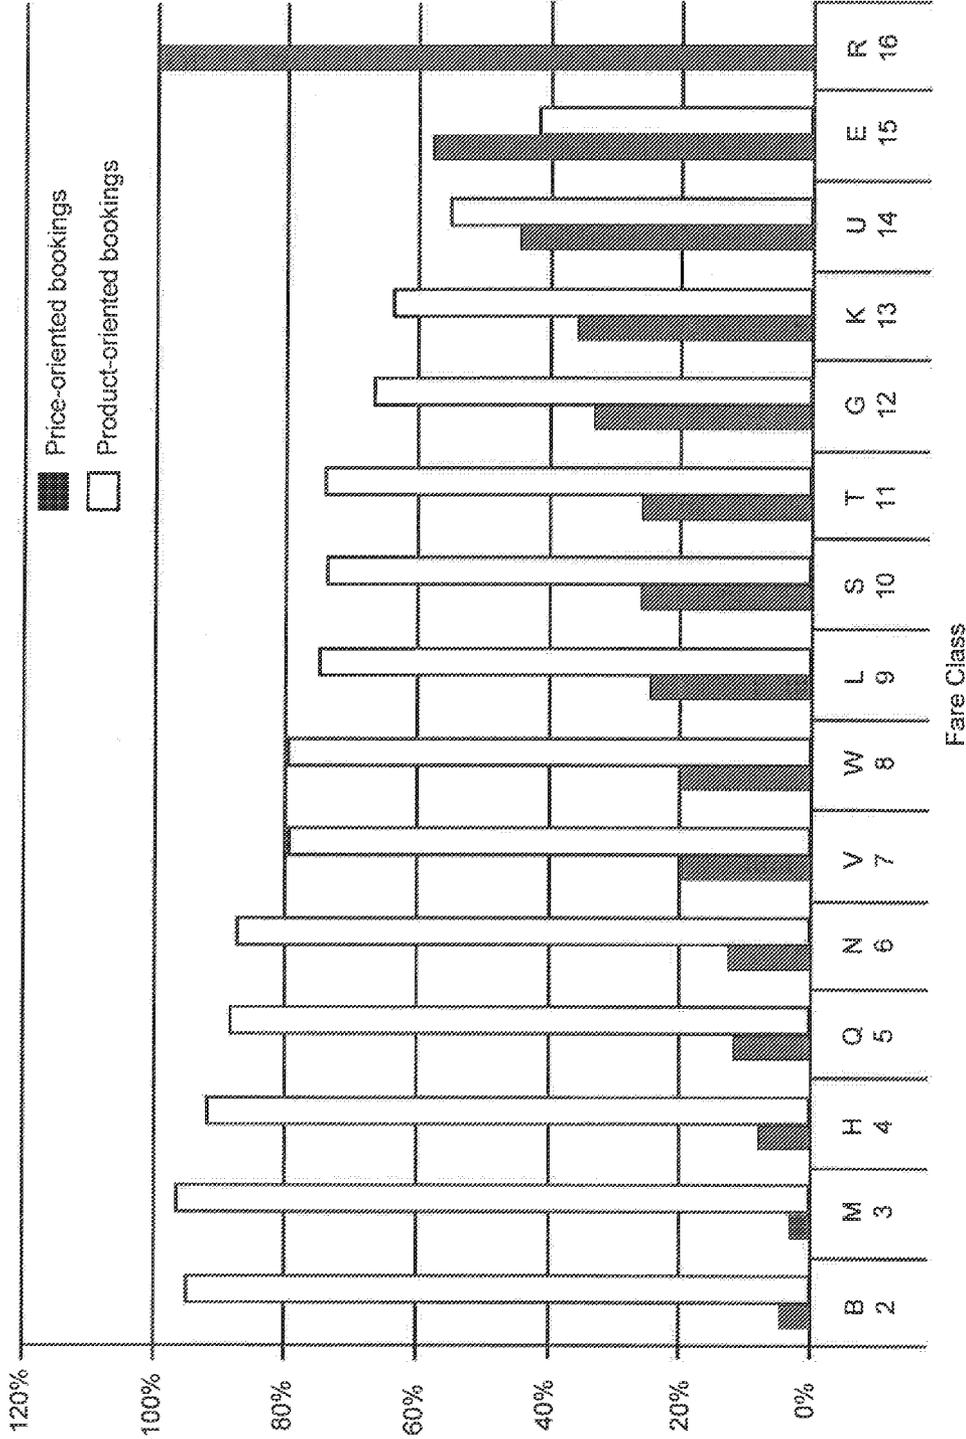


FIG. 4B

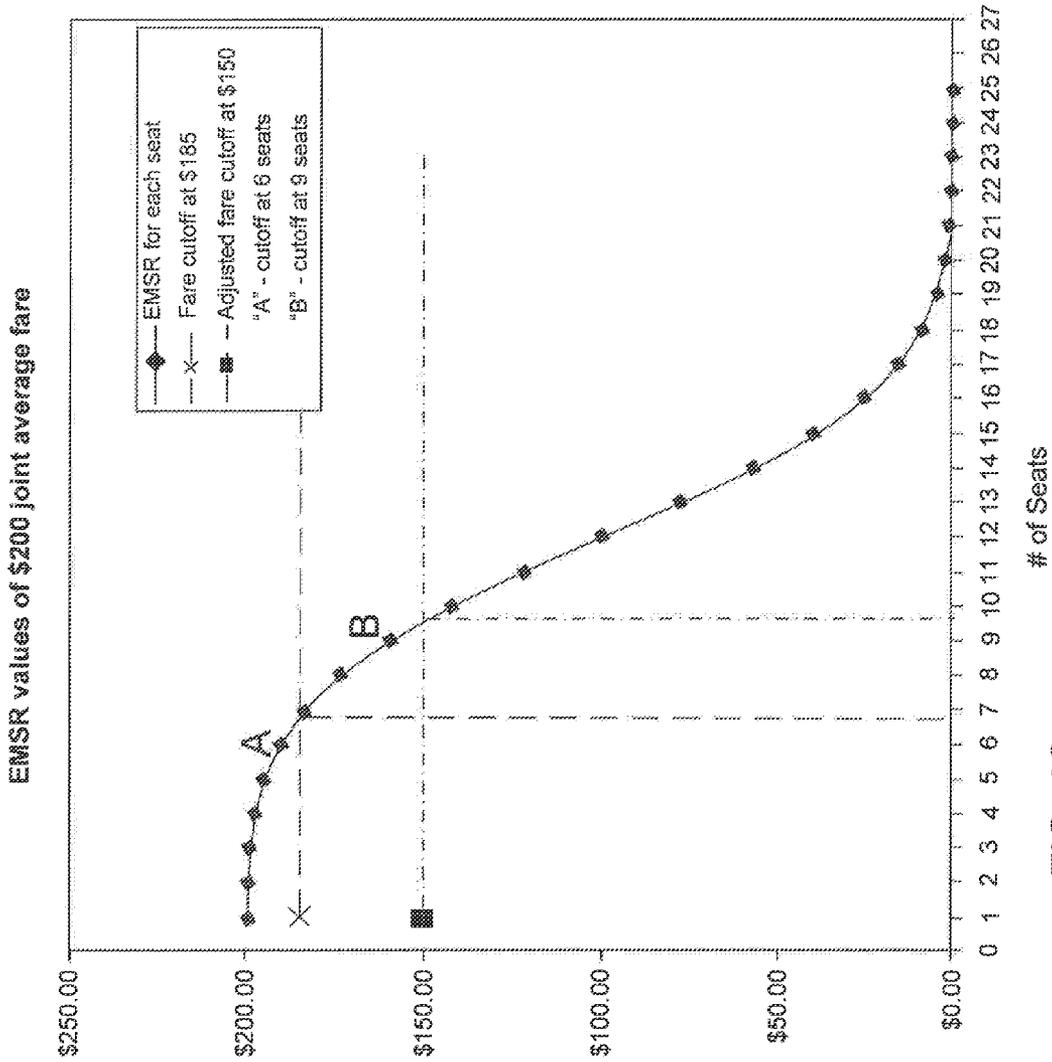


FIG. 4C

500

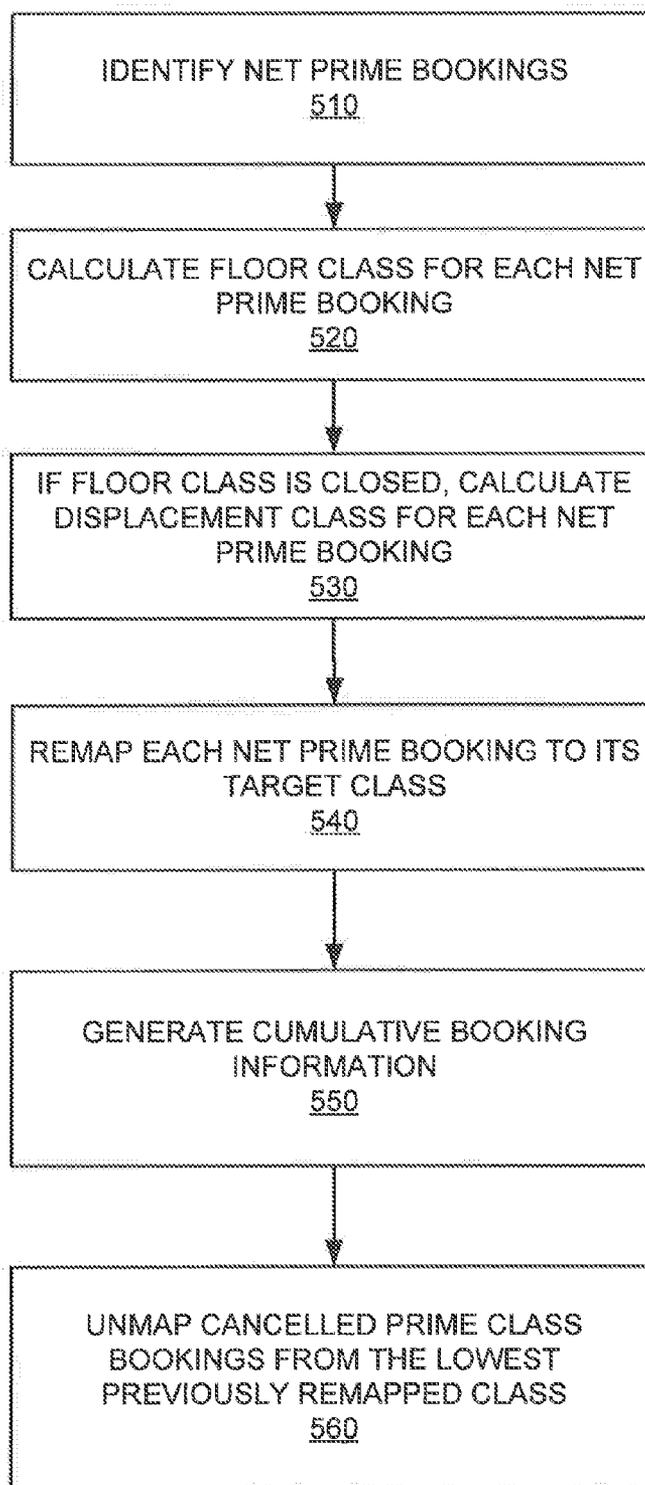


FIG. 5

DEMAND FORECASTING SYSTEMS AND METHODS UTILIZING FARE ADJUSTMENT

TECHNICAL FIELD

[0001] The present disclosure generally relates to forecasting, and more particularly, to analysis methods and tools suitable for use in connection with yield management systems, inventory control systems, revenue management systems, and the like.

BACKGROUND

[0002] The transportation services industry, and particularly the airline industry, is often associated with high costs and varying degrees of profitability. As a result, airlines often seek new sources of income (e.g., ala carte pricing for additional services) and innovative ways to increase revenues and/or reduce costs (e.g., optimizing existing processes). One such method of increasing revenues involves pricing seats on a particular flight in groups or “classes”. Bookings in higher-fare classes are typically given priority in order to increase the total revenue associated with a flight.

[0003] However, in many instances, a passenger may book a reservation at a lower price point than the maximum amount the passenger would be willing to pay for that seat. The difference in price between the higher-priced ticket and the purchased ticket represents foregone revenue for the airline. Accordingly, improved forecasting and analysis methods remain desirable, for example, for use in systems configured to allow airlines to book seats at prices more closely associated with the maximum amount a passenger is willing to pay.

SUMMARY

[0004] In an embodiment, a method comprises accessing, by a processor for adjusting fares, a bookings table containing seat bookings and forecasted demand for a flight. The seat bookings are grouped into a plurality of fare classes, and at least one fare class in the plurality of fare classes is closed and contains no bookings. The method further comprises: determining, by the processor and using the seat bookings and the forecasted demand, the monetary stimulation value of opening the closed class; determining, by the processor and using the seat bookings and the forecasted demand, the monetary dilution penalty of opening the closed class; and comparing, by the processor, the monetary stimulation value to the monetary dilution penalty to generate booking instructions.

[0005] In another embodiment, a non-transitory computer-readable storage medium has computer-executable instructions stored thereon that, in response to execution by a processor for adjusting fares, causes the processor to perform operations comprising: accessing, by the processor, a bookings table containing seat bookings and forecasted demand for a flight, wherein the seat bookings are grouped into a plurality of fare classes, and wherein at least one fare class in the plurality of fare classes is closed and contains no bookings; determining, by the processor and using the seat bookings and the forecasted demand, the monetary stimulation value of opening the closed class; determining, by the processor and using the seat bookings and the forecasted demand, the monetary dilution penalty of opening the closed class; and comparing, by the processor, the monetary stimulation value to the monetary dilution penalty to generate booking instructions.

[0006] The contents of this summary section are provided only as a simplified introduction to the disclosure, and are not intended to be used to limit the scope of the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0007] With reference to the following description, appended claims, and accompanying drawings:

[0008] FIG. 1A is a block diagram illustrating exemplary forecasting system components in accordance with various embodiments;

[0009] FIG. 1B illustrates exemplary approaches for yield management, including reducing spiral down in accordance with various embodiments;

[0010] FIG. 2A illustrates an exemplary unobscuring method in accordance with various embodiments;

[0011] FIG. 2B illustrates an exemplary unobscuring and unconstraining method in accordance with various embodiments;

[0012] FIG. 2C illustrates the result of an exemplary unobscuring method in accordance with various embodiments;

[0013] FIG. 3A illustrates an exemplary unconstraining method in accordance with various embodiments;

[0014] FIG. 3B illustrates the result of an exemplary unconstraining method in accordance with various embodiments;

[0015] FIG. 4A illustrates an exemplary fare adjustment method in accordance with various embodiments;

[0016] FIG. 4B illustrates a chart of exemplary price-oriented bookings and product-oriented bookings in a bookings table in accordance with various embodiments;

[0017] FIG. 4C illustrates seat allocation in a fare class in accordance with an exemplary fare adjustment method, in accordance with various embodiments; and

[0018] FIG. 5 illustrates a method for prime class remapping in accordance with various embodiments.

DETAILED DESCRIPTION

[0019] Principles of the present disclosure can reshape the way organizations forecast, calculate, and/or implement decisions, such as revenue management and/or cost reduction strategies. For example, principles of the present disclosure enable airlines to:

[0020] formulate more accurate answers to the question, “How many (and what type) of bookings would have been realized if a particular fare class (or set of fare classes) were closed?”;

[0021] formulate more accurate answers to the question, “How many (and what type) of bookings would have been realized if a particular class (or set of fare classes) were opened?”;

[0022] realize bookings that are more closely aligned with the highest price a customer is actually willing to pay for a particular seat;

[0023] adjust, at discrete intervals and/or in real-time, demand forecasts for particular fare classes and implement fare class openings/closings in order to increase and/or maximize revenue; stated generally, principles of the present disclosure enable airlines to reduce forecasting errors and account for interdependence of demand between fare classes.

[0024] While the present disclosure discusses airlines and “flights” for purposes of convenience and illustration, one of skill in the art will appreciate that the forecasting methods, systems, and tools disclosed herein are broadly applicable,

for example to any transportation industry, such as buses, cruise ships, passenger trains, and the like.

[0025] Various embodiments of principles of the present disclosure employ forecasting, statistical analysis and/or optimization techniques. For more information regarding such techniques refer to, for example: “The Theory and Practice of Revenue Management” (International Series in Operations Research & Management Science) by Kalyan T. Talluri and Garrett J. van Ryzin; “Using Multivariate Statistics (5th Edition)” by Barbara G. Tabachnick and Linda S. Fidell; and “Introduction to Operations Research” by Friedrich S. Hiller and Gerald J. Lieberman, McGraw-Hill 7th edition, Mar. 22, 2002; the contents of which are each hereby incorporated by reference in their entireties.

[0026] In various embodiments, exemplary forecasting systems include a user interface (“UI”), software modules, logic engines, various databases, interfaces to systems and tools, and/or computer networks. While exemplary forecasting systems may contemplate upgrades or reconfigurations of existing processing systems, changes to existing databases and system tools are not necessarily required by principles of the present disclosure.

[0027] The benefits provided by principles of the present disclosure include, for example, increased revenue, increased forecasting accuracy, lower costs, increased seat utilization, increased customer good will, increased planning and operational efficiency, and increased employee morale. For example, a revenue management organization benefits from improved forecasting accuracy and resulting increased revenue. Customers benefit from booking availability that more closely tracks their willingness to pay, increasing the likelihood of a suitable seat and associated fare being available for them.

[0028] As used herein: a “fare class” or “class” refers to a group of airline seats that are priced similarly to one another.

[0029] “Willingness to pay” refers to the highest fare class a customer will pay for a product, such as a seat on a flight (i.e., analogous to a reserve price).

[0030] “Buy down” refers to a booking being realized in a class that is lower than the customer’s willingness to pay.

[0031] “Spiral down” refers to repeated buy down cycles; when forecasts and seat protects are based on observed bookings rather than willingness to pay, leading to future bookings in lower classes.

[0032] A “constrained class” is a class where bookings are restricted by availability in that class.

[0033] An “obscured class” is a class where bookings are restricted due to availability of the same product in a lower fare class.

[0034] An “active class” or “observed class” is the lowest class with availability and bookings. Lower classes are constrained and higher classes are obscured.

[0035] A “parent class” and “child class” are the fare classes directly above and below the other, respectively.

[0036] A “cohort” is a group of flights having similar characteristics, such that statistical information for some (or all) of the flights in the cohort may be utilized to draw inferences regarding one or more flights in the cohort.

[0037] An “entity” may include any individual, software program, business, organization, government entity, web site, system, hardware, and/or any other entity.

[0038] A “user” may include any entity that interacts with a system and/or participates in a process.

[0039] Turning now to FIG. 1A, in accordance with various embodiments, a user **105** may perform tasks such as requesting, retrieving, receiving, updating, analyzing and/or modifying data. User **105** may also perform task such as initiating, manipulating, interacting with or using a software application, tool, module or hardware, and initiating, receiving or sending a communication. User **105** may interface with Internet server **125** via any communication protocol, device or method discussed herein, known in the art, or later developed. User **105** may be, for example, a member of a revenue management organization, a member of an operations research and systems analysis organization, a downstream system, an upstream system, a third-party system, a system administrator, and/or the like.

[0040] In various embodiments, a system **101** may include a user **105** interfacing with a forecasting system **115** by way of a client **110**. Forecasting system **115** may be a partially or fully integrated system comprised of various subsystems, modules and databases. Client **110** comprises any hardware and/or software suitably configured to facilitate entering, accessing, requesting, retrieving, updating, analyzing and/or modifying data. The data may include operational data (e.g., schedules, resources, routes, operational alerts, weather, etc.), passenger data, cost data, forecasts, historical data, verification data, asset (e.g., airplane) data, inventory (e.g., airplane seat) data, legal/regulatory data, authentication data, demographic data, transaction data, or any other suitable information discussed herein.

[0041] Client **110** includes any device (e.g., a computer), which communicates, in any manner discussed herein, with forecasting system **115** via any network or protocol discussed herein. Browser applications comprise Internet browsing software installed within a computing unit or system to conduct online communications and transactions. These computing units or systems may take the form of personal computers, mobile phones, personal digital assistants, mobile email devices, laptops, notebooks, hand-held computers, portable computers, kiosks, and/or the like. Practitioners will appreciate that client **110** may or may not be in direct contact with forecasting system **115**. For example, client **110** may access the services of forecasting system **115** through another server, which may have a direct or indirect connection to Internet server **125**. Practitioners will further recognize that client **110** may present interfaces associated with a software application (e.g., SAS analytic software) or module that are provided to client **110** via application GUIs or other interfaces and are not necessarily associated with or dependent upon internet browsers or internet specific protocols.

[0042] User **105** may communicate with forecasting system **115** through a firewall **120**, for example to help ensure the integrity of forecasting system **115** components. Internet server **125** may include any hardware and/or software suitably configured to facilitate communications between the client **110** and one or more forecasting system **115** components.

[0043] Firewall **120**, as used herein, may comprise any hardware and/or software suitably configured to protect forecasting system **115** components from users of other networks. Firewall **120** may reside in varying configurations including stateful inspection, proxy based and packet filtering, among others. Firewall **120** may be integrated as software within Internet server **125**, any other system **101** component, or may reside within another computing device or may take the form of a standalone hardware component.

[0044] Authentication server **130** may include any hardware and/or software suitably configured to receive authentication credentials, encrypt and decrypt credentials, authenticate credentials, and/or grant access rights according to pre-defined privileges associated with the credentials. Authentication server **130** may grant varying degrees of application and/or data level access to users based on information stored within authentication database **135** and user database **140**. Application server **142** may include any hardware and/or software suitably configured to serve applications and data to a connected client **110**.

[0045] In accordance with various embodiments, forecasting system **115** is usable to increase and/or maximize revenue, manage inventory strategy, generate inputs to other forecasting systems, and/or the like. Continuing to reference FIG. 1A, forecasting system **115** allows communication with central data repository (CDR) **150**, and with various other databases, tools, UIs and systems (not shown in FIG. 1A). Such systems include, for example, airline scheduling systems, passenger booking and reservations systems, revenue management systems, inventory systems, and/or the like.

[0046] Forecasting system **115** components are interconnected and communicate with one another to allow for a completely integrated forecasting system. In various embodiments, forecasting system **115** formulates demand forecast models and associated revenue consequences at the class level, seat level, and so forth. Airline reservations systems sell inventory based at least in part upon the output of forecasting system **115**.

[0047] In various embodiments, forecasting system **115** modules (e.g., willingness to pay (WTP) system **145**, unobscuring module **146**, unconstraining module **147**, fare adjustment module **148**, prime-class remap module **149**, and other forecasting system **115** modules not shown in FIG. 1A) are software modules configured to enable online functions such as sending and receiving messages, receiving query requests, configuring responses, dynamically configuring user interfaces, requesting data, receiving data, displaying data, executing complex processes, calculations, forecasts, mathematical techniques, workflows and/or algorithms, prompting user **105**, verifying user responses, authenticating the user, initiating forecasting system **115** processes, initiating other software modules, triggering downstream systems and processes, encrypting and decrypting, and/or the like. Additionally, forecasting system **115** modules may include any hardware and/or software suitably configured to receive requests from client **110** via Internet server **125** and application server **142**.

[0048] Forecasting system **115** modules may be further configured to process requests, execute transactions, construct database queries, and/or execute queries against databases within system **101** (e.g., central data repository ("CDR") **150**), external data sources and/or temporary databases. In various embodiments, one or more forecasting system **115** modules may be configured to execute application programming interfaces in order to communicate with a variety of messaging platforms, such as email systems, wireless communications systems, mobile communications systems, multimedia messaging service ("MMS") systems, short messaging service ("SMS") systems, and the like.

[0049] Forecasting system **115** modules may be configured to exchange data with other systems and application modules, for example an airline reservation system. In various embodiments, forecasting system **115** modules may be configured to

interact with other system **101** components to perform complex calculations, retrieve additional data, format data into reports, create XML representations of data, construct markup language documents, construct, define or control UIs, and/or the like. Moreover, forecasting system **115** modules may reside as standalone systems or tools or may be incorporated with the application server **142** or any other forecasting system **115** component as program code. As one of ordinary skill in the art will appreciate, forecasting system **115** modules may be logically or physically divided into various subcomponents, such as a workflow engine configured to evaluate predefined rules and to automate processes.

[0050] In addition to the components described above, forecasting system **115** may further include one or more of the following: a host server or other computing systems including a processor for processing digital data; a memory coupled to the processor for storing digital data; an input digitizer coupled to the processor for inputting digital data; an application program stored in the memory and accessible by the processor for directing processing of digital data by the processor; a display device coupled to the processor and memory for displaying information derived from digital data processed by the processor; a plurality of databases, and/or the like.

[0051] As will be appreciated by one of ordinary skill in the art, one or more system **101** components may be embodied as a customization of an existing system, an add-on product, upgraded software, a stand-alone system (e.g., kiosk), a distributed system, a method, a data processing system, a device for data processing, and/or a computer program product. Accordingly, individual system **101** components may take the form of an entirely software embodiment, an entirely hardware embodiment, or an embodiment combining aspects of both software and hardware. Furthermore, individual system **101** components may take the form of a computer program product on a computer-readable storage medium having computer-readable program code means embodied in the storage medium. Any suitable computer-readable storage medium may be utilized, including magnetic storage devices (e.g., hard disks), optical storage devices, (e.g., DVD-ROM, CD-ROM, etc.), electronic storage devices (e.g., flash memory), and/or the like.

[0052] Client **110** may include an operating system (e.g., Windows, UNIX, Linux, Solaris, MacOS, iOS, Windows Mobile OS, Windows CE, Palm OS, Symbian OS, Blackberry OS, J2ME, etc.) as well as various conventional support software and drivers typically associated with mobile devices and/or computers. Client **110** may be in any environment with access to any network, including both wireless and wired network connections. In various embodiments, access is through a network or the Internet through a commercially available web-browser software package. Client **110** and forecasting system **115** components may be independently, separately or collectively suitably coupled to the network via data links which include, for example, a connection to an Internet Service Provider (ISP) over the local loop as is typically used in connection with standard wireless communications networks and/or methods, such as modem communication, cable modem, satellite networks, ISDN, digital subscriber line (DSL), and/or the like. In various embodiments, any portion of client **110** may be partially or fully connected to a network using a wired ("hard wire") connec-

tion. As those skilled in the art will appreciate, client **110** and/or any of the system components may include wired and/or wireless portions.

[0053] Internet server **125** may be configured to transmit data to client **110** within markup language documents. "Data" may include encompassing information such as commands, messages, transaction requests, queries, files, data for storage, and/or the like in digital or any other form. Internet server **125** may operate as a single entity in a single geographic location or as separate computing components located together or in separate geographic locations. Further, Internet server **125** may provide a suitable web site or other Internet-based graphical user interface, which is accessible by users (such as user **105**). In various embodiments, Microsoft Internet Information Server (IIS), Microsoft Transaction Server (MTS), and Microsoft SQL Server, are used in conjunction with a Microsoft operating system, Microsoft NT web server software, a Microsoft SQL Server database system, and a Microsoft Commerce Server. In various embodiments, the well-known "LAMP" stack (Linux, Apache, MySQL, and PHP/Perl/Python) are used to enable forecasting system **115**. Additionally, components such as Access or Microsoft SQL Server, Oracle, Sybase, Interbase, etc., may be used to provide an Active Data Object (ADO) compliant database management system.

[0054] Like Internet server **125**, application server **142** may communicate with any number of other servers, databases and/or components through any means known in the art. Further, application server **142** may serve as a conduit between client **110** and the various systems and components of forecasting system **115**. Internet server **125** may interface with application server **142** through any means known in the art including a LAN/WAN, for example. Application server **142** may further invoke software modules, such as WTP system **145**, automatically or in response to user **105** requests.

[0055] Any of the communications, inputs, storage, databases or displays discussed herein may be facilitated through a web site having web pages. The term "web page" as it is used herein is not meant to limit the type of documents and applications that may be used to interact with the user. For example, a typical web site may include, in addition to standard HTML documents, various forms, Java applets, JavaScript, active server pages (ASP), common gateway interface scripts (CGI), Flash files or modules, FLEX, ActionScript, extensible markup language (XML), dynamic HTML, cascading style sheets (CSS), helper applications, plug-ins, and/or the like. A server may include a web service that receives a request from a web server, the request including a URL (e.g., <http://yahoo.com/>) and/or an internet protocol ("IP") address. The web server retrieves the appropriate web pages and sends the data or applications for the web pages to the IP address. Web services are applications that are capable of interacting with other applications over a communications means, such as the Internet. Web services are typically based on standards or protocols such as XML, SOAP, WSDL and UDDI. Web services methods are well known in the art, and are covered in many standard texts. See, e.g., Alex Nghiem, *IT Web Services: A Roadmap for the Enterprise* (2003).

[0056] Continuing to reference FIG. 1A, illustrated are databases that are included in various embodiments. An exemplary list of various databases used herein includes: an authentication database **135**, a user database **140**, CDR **150** and/or other databases that aid in the functioning of the system. As practitioners will appreciate, while depicted as sepa-

rate and/or independent entities for the purposes of illustration, databases residing within system **101** may represent multiple hardware, software, database, data structure and networking components. Furthermore, embodiments are not limited to the databases described herein, nor do embodiments necessarily utilize each of the disclosed databases.

[0057] Authentication database **135** may store information used in the authentication process such as, for example, user identifiers, passwords, access privileges, user preferences, user statistics, and the like. User database **140** maintains user information and credentials for forecasting system **115** users (e.g., user **105**).

[0058] In various embodiments, CDR **150** is a data repository that may be configured to store a wide variety of comprehensive data for forecasting system **115**. While depicted as a single logical entity in FIG. 1A, those of skill in the art will appreciate that CDR **150** may, in various embodiments, consist of multiple physical and/or logical data sources. In various embodiments, CDR **150** stores operational data, schedules, resource data, asset data, inventory data, personnel information, routes and route plans, station (e.g., airports or other terminals) data, operational alert data, weather information, passenger data, reservation data, cost data, optimization results, booking class data, forecasts, historical data, verification data, authentication data, demographic data, legal data, regulatory data, transaction data, security profiles, access rules, content analysis rules, audit records, predefined rules, process definitions, financial data, and the like. For example, in various exemplary embodiments a data source or component database of CDR **150** includes information such as flight key (e.g., flight number, flight date, origin, and destination), forecast timeband, days until departure ("Rel Day"), class demand, a class posted indicator (i.e., an indication of whether or not a class was available for sale during a time period), class bookings information, class average fare, active class by Rel Day, seasonal identifiers, sponsor identifiers, and/or the like.

[0059] Any databases discussed herein may include relational, hierarchical, graphical, or object-oriented structure and/or any other database configurations. Common database products that may be used to implement the databases include DB2 by IBM (Armonk, N.Y.), various database products available from Oracle Corporation (Redwood Shores, Calif.), Microsoft Access or Microsoft SQL Server by Microsoft Corporation (Redmond, Wash.), MySQL by MySQL AB (Uppsala, Sweden), or any other suitable database product. Moreover, the databases may be organized in any suitable manner, for example, as data tables or lookup tables. Each record may be a single file, a series of files, a linked series of data fields or any other data structure. Association of certain data may be accomplished through any desired data association technique such as those known or practiced in the art. For example, the association may be accomplished either manually or automatically. Automatic association techniques may include, for example, a database search, a database merge, GREP, AGREP, SQL, using a key field in the tables to speed searches, sequential searches through all the tables and files, sorting records in the file according to a known order to simplify lookup, and/or the like. The association step may be accomplished by a database merge function, for example, using a "key field" in pre-selected databases or data sectors. Various database tuning steps are contemplated to optimize database performance. For example, frequently used files

such as indexes may be placed on separate file systems to reduce In/Out (“I/O”) bottlenecks.

[0060] One skilled in the art will also appreciate that, for security reasons, any databases, systems, devices, servers or other components of system **101** may consist of any combination thereof at a single location or at multiple locations, wherein each database or system includes any of various suitable security features, such as firewalls, access codes, encryption, decryption, compression, decompression, and/or the like.

[0061] The systems and methods may be described herein in terms of functional block components, screen shots, optional selections and various processing steps. It should be appreciated that such functional blocks may be realized by any number of hardware and/or software components configured to perform the specified functions. For example, the system may employ various integrated circuit components, e.g., memory elements, processing elements, logic elements, look-up tables, and the like, which may carry out a variety of functions under the control of one or more microprocessors or other control devices. Similarly, the software elements of the system may be implemented with any programming or scripting language such as C, C++, C#, Java, JavaScript, Flash, ActionScript, FLEX, VBScript, Macromedia Cold Fusion, COBOL, Microsoft Active Server Pages, assembly, PERL, SAS, PHP, awk, Python, Visual Basic, SQL Stored Procedures, PL/SQL, any UNIX shell script, and/or extensible markup language (XML) or the like, with the various algorithms being implemented with any combination of data structures, objects, processes, routines or other programming elements. Further, it should be noted that the system may employ any number of conventional techniques for data transmission, signaling, data processing, network control, and the like. Still further, the system may be used to detect or prevent security issues with a client-side scripting language, such as JavaScript, VBScript or the like.

[0062] Software elements may be loaded onto a general purpose computer, special purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions that execute on the computer or other programmable data processing means for implementing the functions specified in the flowchart block or blocks. These computer program instructions may also be stored in a computer-readable memory that can direct a computer or other programmable data processing apparatus to function in a particular manner, such that the instructions stored in the computer-readable memory produce an article of manufacture including instruction means which implement the function specified herein or in flowchart block or blocks. The computer program instructions may also be loaded onto a computer or other programmable data processing apparatus to cause a series of operational steps to be performed on the computer or other programmable apparatus to produce a computer-implemented process such that the instructions which execute on the computer or other programmable apparatus provide steps for implementing the functions specified in the flowchart block or blocks.

[0063] Accordingly, functional blocks of the block diagrams and flowchart illustrations support combinations of means for performing the specified functions, combinations of steps for performing the specified functions, and program instruction means for performing the specified functions. It will also be understood that each functional block of the block diagrams and flowchart illustrations, and combinations of

functional blocks in the block diagrams and flowchart illustrations, can be implemented by either special purpose hardware-based computer systems which perform the specified functions or steps, or suitable combinations of special purpose hardware and computer instructions. Further, illustrations of the process flows and the descriptions thereof may make reference to user windows, web pages, web sites, web forms, prompts, etc. Practitioners will appreciate that the illustrated steps described herein may comprise any number of configurations including the use of windows, web pages, web forms, popup windows, prompts and/or the like. It should be further appreciated that the multiple steps as illustrated and described may be combined into single web pages and/or windows but have been expanded for the sake of simplicity. In other cases, steps illustrated and described as single process steps may be separated into multiple web pages and/or windows but have been combined for simplicity.

[0064] With continued reference to FIG. 1A, in various embodiments, user **105** logs onto an application (e.g., a module) and Internet server **125** may invoke an application server **142**. Application server **142** invokes logic in the forecasting system **115** modules by passing parameters relating to user's **105** requests for data. Forecasting system **115** manages requests for data from forecasting system **115** modules and/or communicates with system **101** components. Transmissions between user **105** and Internet server **125** may pass through a firewall **120** to help ensure the integrity of forecasting system **115** components. Practitioners will appreciate that exemplary embodiments may incorporate any number of security schemes or none at all. In various embodiments, Internet server **125** receives requests from client **110** and interacts with various other system **101** components to perform tasks related to requests from client **110**.

[0065] Internet server **125** may invoke an authentication server **130** to verify the identity of user **105** and assign roles, access rights and/or permissions to user **105**. In order to control access to the application server **142** or any other component of forecasting system **115**, Internet server **125** may invoke an authentication server **130** in response to user **105** submissions of authentication credentials received at Internet server **125**. In response to a request to access system **101** being received from Internet server **125**, Internet server **125** determines if authentication is required and transmits a prompt to client **110**. User **105** enters authentication data at client **110**, which transmits the authentication data to Internet server **125**. Internet server **125** passes the authentication data to authentication server **130** which queries the user database **140** for corresponding credentials. In response to user **105** being authenticated, user **105** may access various applications and their corresponding data sources.

[0066] With reference now to FIGS. 1A and 1B, various prior approaches to airline seat demand forecasting have suffered from various shortcomings, for example susceptibility to spiral down. Prior demand forecasts were typically based on actual seat bookings, leading to repeated cycles of fare deterioration as additional groups of lower priced seats were released for sale based on inaccurate demand estimates. Additionally, prior demand forecasts typically assumed, inaccurately, that demand for a particular fare class was independent of demand for other fare classes.

[0067] In contrast, principles of the present disclosure contemplate approaches to demand forecasting, wherein interrelationships among demand for fare classes are taken into account. For example, interrelationships among demand for

fare classes may arise anytime that there is more than one valid fare for a particular seat on a flight. In various embodiments, airline seat demand forecasting systems and methods are configured to consider that a particular fare class, while open, may be effectively closed due to seat availability in a lower fare class or classes. Such open fare classes may be referred to as “obscured”, as demand for such fare classes is not readily apparent. Principles of the present disclosure may be utilized to “unobscure” (or otherwise improve) demand forecasts for such fare classes.

[0068] In various embodiments, unobscuring module 146 is configured to generate and/or modify an improved demand forecast for a fare class or set of fare classes. Moreover, unobscuring module 146 may be configured to revise and/or modify historical booking information to represent unobscured booking information. Stated generally, unobscuring module 146 is configured to provide an answer to the question: “For a particular fare class F_i , how many bookings would have been taken if the fare classes below F_i were closed?”

[0069] Unobscuring module 146 may utilize historical data regarding seat bookings for a prior period of time, for example the prior 52 weeks of booking information. Moreover, any suitable amount of historical data may be utilized. Unobscuring module 146 may also utilize information such as flight key (e.g., flight number, flight date, origin, and destination), forecast timeband, days until departure (“Rel Day”), a class posted indicator, class bookings information, seasonal identifiers, sponsor identifiers, and/or the like.

[0070] In various embodiments, unobscuring module 146 is configured to generate predicted booking information that differs from actual observed bookings. In various embodiments, unobscuring module 146 generates demand forecasts directly; in other exemplary embodiments, predicted booking information generated by unobscuring module 146 may be delivered to other electronic systems (for example, Passenger Revenue Optimization System or “PROS” offered by PROS Holdings, Inc.) in lieu of actual observed booking information, in order to generate improved demand forecasts.

[0071] In various embodiments, unobscuring module 146 processes net individual seat bookings that occurred between a first point in time (which may be a “data collection point” or DCP) and a second point in time (which may be, but is not necessarily, a second DCP). Table 1A illustrates an exemplary snapshot of total bookings for a particular example Flight 244, scheduled for travel from Detroit (DTW) to Phoenix (PHX). In Table 1A, total bookings are represented as of a particular DCP (for example, at example DCP 19, which is a point in time 8 days prior to the scheduled departure of Flight 244). In Table 1A, total bookings are listed by fare class in descending price order.

TABLE 1A

Total bookings for Flight 244 as of DCP 19	
Fare Class	Total Bookings to Date
Q	1
N	1
V	8
W	15
L	9
S	7
T	19
G	7

TABLE 1A-continued

Total bookings for Flight 244 as of DCP 19	
Fare Class	Total Bookings to Date
K	15
U	7
E	5
R	0

[0072] Continuing this example, in the time period between the current DCP 19 and the prior DCP 18, the following new bookings for Flight 244 occurred:

TABLE 1B

New Bookings for Flight 244 Occurring Between DCP 18 and DCP 19	
Fare Class	New Bookings
Q	0
N	0
V	0
W	3
L	0
S	0
T	1
G	0
K	4
U	0
E	0
R	0

[0073] To facilitate improved demand forecasting approaches and/or to facilitate interactions with other components of system 101 (for example, external systems and databases 160), the booking information in Tables 1A and/or 1B may also be considered in “top-down” form, wherein the booking number associated with each fare class represents the sum of bookings in that fare class as well as all higher fare classes. Converting the new booking information in Table 1B to top-down form, Table 1C is obtained:

TABLE 1C

Fare Class	New Bookings	New Bookings (Top-Down)
Q	0	0
N	0	0
V	0	0
W	3	3
L	0	3
S	0	3
T	1	4
G	0	4
K	4	8
U	0	8
E	0	8
R	0	8

[0074] With the example data formatted as illustrated in Table 1C, a baseline assumption about bookings for this example Flight 244 can be inferred, namely: if, during the current period (i.e., between DCP 18 and DCP 19), fare classes U through R were not open for bookings, Flight 244 would have received at least 8 bookings (i.e., a number equal to the actual bookings). However, if fare class K was not open

for bookings, Flight 244 might still have received the four bookings taken in fare class K (i.e., in fare class G or another higher class); however, it is also possible that only three, or two, or one, or zero of the observed K class bookings would have occurred. Because of this uncertainty, the actual demand for bookings in fare class G may be considered to be obscured. All fare classes above class G may also be considered obscured.

[0075] With reference now to FIGS. 2A and 2B, in accordance with various embodiments, unobscuring module 146 is configured to implement a method 200 for unobscuring demand. In method 200, unobscuring module 146 may operate class-by-class, for example from the lowest price class through the highest price class, in order to provide a revised demand forecast and/or unobscured booking information useable to reduce and/or minimize forecasting errors associated with obscured demand.

[0076] Unobscuring module 146 and/or method 200 may be utilized on any suitable time schedule, as desired. For example, unobscuring module 146 may be operable on a weekly basis, on a daily basis, on an hourly basis, on a “triggered” basis (for example, each time a trigger event occurs, such as a booking or series of bookings), after collection of data associated with a DCP, and/or the like.

[0077] In various embodiments, and with reference to FIG. 2A, in method 200 booking information for a particular flight is converted into top-down format (step 210). Beginning with the lowest fare class as the “current class”, the mean and standard deviation of top-down bookings in the current class is computed (step 220). In various exemplary embodiments, statistical information from the cohort to which the particular flight belongs is used to generate the mean and standard deviation of top-down bookings in the current class. A normal distribution may initially be assumed; however, other suitable distributions may also be utilized. If the current class is obscured, a suitable estimation algorithm, for example an expectation maximization (EM) algorithm 235 (not shown in the figures), is applied in an iterative fashion to calculate an unobscured demand for the current class (step 230). In various embodiments, EM algorithm 235 comprises an algorithm identical or substantially equivalent to:

$$E(X | a < X < b) = \mu + \frac{\phi\left(\frac{a-\mu}{\sigma}\right) - \phi\left(\frac{b-\mu}{\sigma}\right)}{\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)}\sigma$$

[0078] In an exemplary embodiment, EM algorithm 235 may be written in a non-normalized form as: $\text{expBkgs} = \text{mean} + (\text{stdev} * \text{stdev}) * (\text{pdfLower} - \text{pdfUpper}) / (\text{cdfUpper} - \text{cdfLower})$; where

[0079] expBkgs is the expectation of bookings in the current class; i.e. the unobscured bookings;

upperBound is the upper bound of the unobscured bookings in the current class, which may be equivalent to the unobscured bookings in the child class;

lowerBound is the lower bound of the unobscured bookings in the current class, which may be equivalent to the actual top-down cumulative bookings taken in the current class;

mean : the mean of the bookings (unobscured, unconstrained, and/or unobscured and unconstrained) for the current class in the appropriate cohort to which the current flight belongs;

stdev : the standard deviation of the bookings (unobscured, unconstrained, and/or unobscured and unconstrained) for the current class in the appropriate cohort to which the current flight belongs;

cdfUpper is the cumulative distribution function of the upper bound;

cdfLower is the cumulative distribution function of the lower bound;

pdfUpper is the probability density function of the upper bound; and

pdfLower is the probability density function of the lower bound.

[0080] In various embodiments, in method 200 the unobscured demand for the current class is used as the upper bound for unobscured demand for the parent class. The process is repeated for each parent class (step 240) until the unobscured demand for the highest fare class is obtained.

[0081] Because airline seats are booked on an integer basis (i.e., a partial seat booking is not possible), the unobscured demand values may be rounded or otherwise converted into discrete values (step 250); however, it will be appreciated that (i) when unobscuring module 146 is utilized in connection with other systems configured to accept fractional demand values, and/or (ii) when unobscuring module 146 is utilized in connection with demand forecasting associated with products and/or services amendable to fractional demand modeling, step 250 may be omitted. An updated demand forecast and/or unobscured booking information is generated (step 260).

[0082] In various exemplary embodiments, statistical information from a cohort or cohorts to which the current flight belongs may be utilized in unobscuring module 146 and/or method 200. A cohort may comprise any suitable number of flights; moreover, any suitable amount of historical data for the cohort may be utilized. In various exemplary embodiments, a cohort comprises between about 25 flights and about 75 flights. In one exemplary embodiment, a cohort comprises 50 flights. In various exemplary embodiments, between one month of historical information and two years of historical information (or 3+ years of historical information, for seasonal effects) may be utilized to conduct statistical assessments of the cohort. In one exemplary embodiment, 52 weeks of historical information are utilized to make a statistical assessment of the cohort.

[0083] With continued reference to FIG. 2B, method 200 may be applied to the exemplary booking data for Flight 244 provided in Table 1C. For each fare class, it is known that the minimum demand in that fare class is the top-down value of bookings in that class. Accordingly, with respect to the lowest fare class, R class, the lower bound on demand is the observed topdown bookings (i.e., 8). The upper bound is based on the class below class R (i.e., the child class). However, because class R is the lowest fare class, it has no child class, so the expected demand for R class, DEM_R , is bounded as follows: $8 \leq \text{DEM}_R \leq 8$. Moving up to the parent class of class R, class E, DEM_R becomes the new upper bound for DEM_E . Therefore, $8 \leq \text{DEM}_E \leq 8$. Classes U and K follow in similar fashion.

[0084] In this example, G class is the first class for which demand is considered to be obscured. Here, $4 \leq \text{DEM}_G \leq 8$. The lower bound, 4, is the observed topdown bookings in class G; the upper bound, 8, is DEM_K , the expected demand in K class (the child class of class G). An EM algorithm may be utilized to solve for DEM_G , for example as illustrated in FIG. 2C.

[0085] Using an exemplary EM algorithm, for example an iterative algorithm applied until a convergent solution is

reached, and using appropriate information from the cohort of the example flight, values for G class are calculated as follows: Mean=8.48, Std. Dev.=3.32, and $DEM_G=6.28$. Stated differently, when applied to our exemplary data for example Flight 244, method **200** predicts that in the current period, 6.28 persons would have booked in G class if K class seating had not been available. Repeating the process for each remaining class, using information from the cohort as appropriate, we obtain:

[0086] $4 \leq DEM_T \leq 6.28$; Mean=8.09, Std. Dev.=3.11, and $DEM_T=5.27$.

[0087] $3 \leq DEM_S \leq 5.27$; Mean=6.58, Std. Dev.=2.84, and $DEM_S=4.26$.

[0088] $3 \leq DEM_L \leq 4.46$; Mean=3.46, Std. Dev.=1.33, and $DEM_L=3.62$.

[0089] $3 \leq DEM_W \leq 3.62$; Mean=3.13, Std. Dev.=1.54, and $DEM_W=3.31$.

[0090] $0 \leq DEM_V \leq 3.31$; Mean=0.67, Std. Dev.=0.30, and $DEM_V=0.67$.

[0091] $0 \leq DEM_N \leq 0.67$; Mean=0.09, Std. Dev.=0.03, and $DEM_N=0.10$.

[0092] At this point, it will be appreciated that while the process may continue for each higher fare class (i.e., class Q), the process may also be terminated at this point, as the expected demand for all higher fare classes is approximately zero.

[0093] The expected demand values generated above for each fare class may be converted into integers, for example via simple rounding. Other suitable integer conversion approaches may be utilized. The results of simple rounding are shown in Table 1D:

TABLE 1D

Fare Class	Expected Demand (Top-Down)	Expected Demand, Rounded (Top-Down)
Q	N/A	0
N	0.10	0
V	0.67	1
W	3.31	3
L	3.62	4
S	4.26	4
T	5.27	5
G	6.28	6
K	8.00	8
U	8.00	8
E	8.00	8
R	8.00	8

[0094] To eliminate double counting, the expected demand values may be converted out of top-down format into a discrete format. Additionally, each booking represented in the expected demand may be allocated to any suitable fare class; however, in various embodiments each booking is allocated to the highest fare class that can claim it, as shown in Table 1E:

TABLE 1E

Fare Class	Expected Demand, Rounded (Top-Down)	Expected Demand, Discrete
Q	0	0
N	0	0
V	1	1
W	3	2

TABLE 1E-continued

Fare Class	Expected Demand, Rounded (Top-Down)	Expected Demand, Discrete
L	4	1
S	4	0
T	5	1
G	6	1
K	8	2
U	8	0
E	8	0
R	8	0

[0095] Continuing with the example data at DCP **19** for exemplary Flight 244, the expected or “unobserved” demand for all fare classes may be compared to the actual bookings, as shown in Table 1F:

TABLE 1F

Fare Class	New Bookings (Actual)	Expected Demand (Discrete)	Difference, Actual to Expected
Q	0	0	0
N	0	0	0
V	0	1	+1
W	3	2	-1
L	0	1	+1
S	0	0	0
T	1	1	0
G	0	1	+1
K	4	2	-2
U	0	0	0
E	0	0	0
R	0	0	0

[0096] In various embodiments, the expected demand information output by method **200** and/or unobscuring module **146** may be combined with previously calculated expected demand information (for example, expected demand information calculated for a prior DCP and/or multiple prior DCPs) and/or with prior actual booking information, and the combined demand information may be utilized as an input into an external forecasting system. It will be appreciated that in instances when a flight is only partially booked (i.e., when many seats on a flight remain open), the output of unobscuring module **146** may result in only minor revisions and/or updates to a demand forecast, a seat protection strategy, and/or the like; however, when supply and demand levels are close to one another (or when demand exceeds supply), the output of unobscuring module **146** may result in significant changes to a demand forecast, a seat protection strategy, and/or the like.

[0097] Unobscuring module **146** (and/or method **200**) may also be configured to utilize weighting factors (for example, seasonal adjustments, and/or the like) to further refine demand estimates. Additionally, booking information for holidays or other unusual dates may be considered to be outliers, and associated booking information may be ignored and/or minimally weighted in unobscuring module **146** and/or method **200**.

[0098] In additional exemplary embodiments, unobscuring module **146** may take additional passenger information factors into account, for example passenger name record (“PNR”) characteristics. Based at least in part on PNR characteristics, unobscuring module **146** may estimate how likely

a particular passenger is to “buy down” when booking a seat, and the likelihood of buy-down may be utilized to refine and/or modify the output of unobscuring module 146.

[0099] Via use of unobscuring module 146 and/or method 200, demand may be more accurately forecasted. Consequently, improved fare class opening/closing decisions may be made.

[0100] In addition to utilizing unobscuring module 146, forecasting system 115 may be configured to utilize unconstraining module 147 (and/or method 300) to create, revise, and/or modify demand forecasts and/or to generate unconstrained demand information. Stated generally, unconstraining module 147 is configured to provide an answer to the question: “How many bookings would have been taken if a particular fare class F_{NEW} was opened?”

[0101] Often, a fare class F_{NEW} is opened partway through the current period, for example the period of time between a first DCP and a second DCP (or other interval of time over which statistical assessment is desirable). Stated another way, often a newly opened fare class F_{NEW} was open for bookings for only a portion of a time period for which booking information is available. It will be appreciated that therefore, there is at least a possibility that additional bookings would have been realized in class F_{NEW} if class F_{NEW} was open for bookings for the entirety of the current period. Because a potential exists for denied bookings in class F_{NEW} , class F_{NEW} may be considered to be “constrained”. Accordingly, the number of bookings that might have been realized in class F_{NEW} is not readily apparent.

[0102] In various exemplary embodiments, and with reference now to FIGS. 1A and 3A, to address this issue unconstraining module 147 is configured to utilize method 300 to obtain unconstrained demand information, for example unconstrained demand information associated with a particular flight. In various embodiments, unconstraining module 147 is configured to generate predicted booking information that differs from actual observed bookings. In certain exemplary embodiments, unconstraining module 147 generates demand forecasts directly; in other exemplary embodiments, predicted booking information generated by unconstraining module 147 may be delivered to other electronic systems (for example, PROS) in lieu of actual observed booking information, in order to generate improved demand forecasts.

[0103] Unconstraining module 147 and/or method 300 may be utilized on any suitable time schedule, as desired. For example, unconstraining module 147 may be operable on a weekly basis, on a daily basis, on an hourly basis, on a “triggered” basis (for example, each time a trigger event occurs, such as a booking or series of bookings), after collection of data associated with a DCP, and/or the like. In various embodiments, unconstraining module 147 is operable subsequent to operation of unobscuring module 146.

[0104] With reference now to FIG. 3A, in various embodiments, in method 300 an unobscured booking table containing unobscured demand information is utilized as an input (step 310). Constrained fare classes are identified (step 320). For each constrained fare class, the corresponding unobscured demand information is converted to unconstrained demand information (step 330), for example by utilizing statistical information from the cohort to which the current flight belongs. In various exemplary embodiments, the statistical information from the cohort comprises an evaluation of a data set of truncated histories (i.e., historical information for previous constrained and unconstrained classes in the cohort);

stated another way, the statistical information from the cohort comprises a comparison of the current constrained fare class to overall statistical information for the current fare class in the cohort as a whole.

[0105] Because airline seats are booked on an integer basis (i.e., a partial seat booking is not possible), the unconstrained demand information may be rounded or otherwise converted into discrete values (step 340); however, it will be appreciated that (i) when unconstraining module 147 is utilized in connection with other systems configured to accept fractional demand values, and/or (ii) when unconstraining module 147 is utilized in connection with demand forecasting associated with products and/or services amendable to fractional demand modeling, step 340 may be omitted. An updated demand forecast and/or unconstrained booking table is generated (step 350).

[0106] Unconstraining module 147 (and/or method 300) may also be configured to utilize weighting factors (for example, seasonal adjustments, and/or the like) to further refine demand estimates. Additionally, booking information for holidays or other unusual dates may be considered to be outliers, and associated booking information may be ignored and/or minimally weighted in unobscuring module 147 and/or method 300.

[0107] Principles of operation of unconstraining module 147 and/or method 300 are illustrated in the example bookings shown in Table 3A. Table 3A illustrates new bookings for an example Flight 344, which received 10 new bookings over a particular interval, for example the time between a first DCP and a second DCP:

TABLE 3A

New Bookings for Flight 344 in the Current Period		
Fare Class	New Bookings	New Bookings (Top-Down)
S	0	0
T	1	1
G	3	4
<u>K</u>	<u>4</u>	8
<u>U</u>	<u>2</u>	10
<u>E</u>	<u>0</u>	10
<u>R</u>	<u>0</u>	10

[0108] In the example illustrated in Table 3A, fare classes K, U, E, and R were opened for bookings partway through the current period (i.e., classes K, U, E, R, were closed for bookings for at least part of the current period). Classes K, U, E, R, are shown as underlined in Table 3A to illustrate that these fare classes were opened during the current period.

[0109] For this example, it can be seen that if classes U, E, and R were open for the entire current period, at least 10 bookings would have occurred. Additionally, if classes U, E, and R, were not opened, the two bookings that occurred in class U might have occurred in class K (however, it is also possible that only one, or none, might have occurred). Accordingly, class K may be considered to be obscured as discussed hereinabove. Additionally, because class K was open for bookings for only a portion of the current period, class K is also considered to be “constrained”.

[0110] In accordance with principles of the present disclosure, unobscuring module 146 may be utilized to determine the unobscured demand for an obscured fare class or classes in this example, such as unobscured demand DEM_K for class

K. In accordance with unobscuring principles disclosed hereinabove, it can be seen that $8 \leq \text{DEM}_K \leq 10$. Via use of unobscuring module 146, DEM_K is determined to be 9.05.

[0111] Because example class K was opened for bookings during the current period, class K is constrained. Unobscured demand DEM_K (in fractional or integer form) is utilized as an input to unconstraining module 147 and/or method 300 to obtain unconstrained and unobscured demand ${}_{\nu}\text{DEM}_K$ for class K. Via operation of unconstraining module 147, ${}_{\nu}\text{DEM}_K$ is determined to be 10.66, for example as illustrated in FIG. 3B. Stated another way, class K would have been expected to receive 10.66 bookings if class K had been open for bookings for the entire current period. A similar approach may be employed to determine ${}_{\nu}\text{DEM}_L$, and similar demand forecasting inferences may be drawn therefrom.

[0112] In various embodiments, through use of unobscuring module 146 and/or unconstraining module 147, airlines may implement revisions to their booking strategies. In many instances, airlines may take more bookings on marginal flights and/or lower revenue seats than before. Stated differently, after assessing the improved estimated demand information provided and/or facilitated by unobscuring module 146 and/or unconstraining module 147, airlines may protect fewer upper-class seats, and more bookings in lower classes are thus made. However, the resulting bookings are typically made in fare classes that are higher than would have resulted if forecasts were based on historical demand alone and/or based on the presumption of fare class demand independence. This is because, after assessment of the improved estimated demand information, buy-down risk in higher classes is better known. Additional seats are thus booked in a manner that is revenue-positive for the airline, as (i) seats previously left empty are occupied, and (ii) buy-down is reduced.

[0113] In addition to unobscuring principles and/or unconstraining principles as disclosed hereinabove, forecasting system 115 may be configured to utilize fare adjustment module 148. Fare adjustment module 148 is configured to provide insight into the demand relationships between one or more fare classes. Fare adjustment module 148 can thus be utilized to determine under what conditions opening a new, lower fare class is a revenue positive decision.

[0114] Prior demand forecasting approaches often assumed independence of demand; in other words, prior approaches often assumed 0% buy down. In contrast, principles of the present disclosure contemplate demand forecasting approaches and/or fare adjustment approaches that recognize that certain passengers may be “product oriented” (i.e., passengers who will purchase in a higher fare class than the lowest available), while certain other passengers may be “price oriented” (i.e., passengers who will buy down, if given the opportunity to do so).

[0115] Because not all passengers are product oriented, and not all passengers are price oriented, opening a new, lower fare class F_L has at least 2 effects. First, “stimulation”. Stimulation refers to revenue arising from and/or represented by new bookings that would not have occurred if class F_L had not been opened. Second, a “dilution penalty” or simply “dilution”. Dilution refers to the negative revenue consequences arising from and/or represented by bookings that were bought down in class F_L , but which would have been purchased in a higher class if class F_L was not open. Stated generally, fare adjustment module 148 may be utilized to evaluate if, for a particular set of circumstances, the stimulation of opening

class F_L exceeds the dilution arising from opening of class F_L —and thus, whether opening class F_L is a revenue positive decision.

[0116] Prior approaches often considered only the effects of stimulation, and did not consider dilution; accordingly, prior approaches often led to class opening decisions that were revenue negative.

[0117] It will be appreciated that the ratio of product-oriented passengers (or product-oriented demand) to price-oriented passengers (or price-oriented demand) has a significant effect on fare adjustments. In fare adjustment module 148, “buy-down percent” may be considered to be the percentage of passengers for which demand is price-oriented. Fare adjustment module 148 may identify, for a particular set of bookings, the associated demand orientation.

[0118] In various embodiments, fare adjustment module 148 is configured to consider certain bookings (for example, bookings that were revised and/or restated as a result of operation of unobscuring module 146) to be price-oriented bookings, as these bookings have already demonstrated buy-down. Certain other bookings may be determined to be price-oriented, for example based on customer purchase history or other historical data. Bookings which are not determined to be price oriented may be considered to be product-oriented. An illustration of exemplary ratios of price-oriented demand vs. product-oriented demand for a series of fare classes on a particular flight is illustrated in FIG. 4B. Stated generally, the proportion of demand for a fare class that is product-oriented demand tends to increase as the expense of the fare class increases; conversely, the proportion of demand for a fare class that is price-oriented demand tends to increase as the expense decreases.

[0119] In various embodiments, fare adjustment module 148 is configured to utilize fare adjustment method 400. With reference now to FIG. 4A, in one embodiment fare adjustment method 400 comprises determining a buy-down percentage (step 410), determining a marginal fare (step 420), and determining an adjusted fare (step 430).

[0120] In fare adjustment method 400, using restated bookings information from unobscuring module 146, a buy-down percentage BD % may be calculated as follows: the ratio of historical bookings in classes higher than the current class that were booked in the lowest available class, or were moved “up” to a higher fare class via unobscuring, to the total number of historical bookings in classes higher than the current class. It will be appreciated that a buy-down percentage may be calculated for a flight or group of flights, for example based on one or more common characteristics of a group of flights. Moreover, BD % may be considered to be the percent of bookings that are price oriented bookings.

[0121] In fare adjustment method 400, a marginal fare MF may be determined as follows: $\text{MF} = (\text{average fare} * \text{probability}(\text{seat purchased by current class})) + (\text{LOSS} * \text{probability}(\text{seat purchased by higher class}))$, where:

[0122] LOSS is the difference in revenue per booking between the current class and the parent class (i.e., the average fare of current class minus the average fare of the parent class);

[0123] $\text{probability}(\text{seat purchased by current class}) = \frac{\text{demand}(\text{current class})}{\text{demand}(\text{current class and all higher classes})}$; and

[0124] $\text{probability}(\text{seat purchased by higher class}) = 1 - \text{probability}(\text{seat purchased by current class})$; the

marginal fare calculation may utilize forward looking demand for a flight or group of flights.

[0125] In fare adjustment method 400, determining an adjusted fare may be implemented as follows: an adjusted fare AdjF for a particular class may be equal to $(BD \% * MF) + (1 - BD \%) * AF$, where AF is the average fare in the class. It will be appreciated that AdjF may be obtained based at least in part on a weighted average of average fare and marginal fare based on buy-down.

[0126] In various embodiments, fare adjustment module 148 makes an initial assessment configured to assume 100% buy down. In other words, if a lower class is opened, everyone from higher classes will buy down to D. This is illustrated by example in Table 4A. In Table 4A, demand for seats on a hypothetical flight having 125 seats is assumed to be deterministic, as shown:

TABLE 4A

Fare Class	Average Fare	Demand	Cumulative/ Top-Down Demand	Total Revenue	Incremental Revenue Arising from Having the Class Open
A	\$400	10	10	\$ 4,000	\$4,000
B	\$300	25	35	\$10,500	\$6,500
C	\$200	40	75	\$15,000	\$4,500
D	\$100	50	125	\$12,500	(\$2,500)

[0127] When modeled in this manner, it is apparent that having class D open is not a revenue positive decision; the true value of each booking in class D (i.e., the marginal fare) is equal to $(\$2500)/125 = (\$20)$, a loss of twenty dollars.

[0128] In accordance with various embodiments, fare adjustment module 148 is configured to utilize expected marginal seat revenue (EMSR) information, for example EMSR

[0131] By way of example, fare adjustment module 148 may consider a hypothetical situation wherein 1 seat protect must be allocated. In this hypothetical situation, the seat protect is in B class, which is determined to have a 20% chance of booking, and a \$900 average fare. The EMSR for this seat is $\$900 * 0.2 = \180 , which, in this example, is the highest EMSR remaining for any open seat in any fare class.

[0132] When fare adjustment module 148 employs a partitioned approach, the seat protect is assigned to B class, which is the class used in the EMSR calculation. An advantage of the partitioned approach is it holds out for the potential higher value seat booking, thus combatting spiral down. However, because the chance of booking the seat is somewhat low (i.e., only 20%), the degree to which spiral down is reduced may be small. Moreover, a potential downside of the partitioned approach is the seat protect in B class will cause potential bookings in lower classes to be rejected, even if those potential bookings have a higher actual value than the EMSR of \$180 for this B class seat.

[0133] Continuing with this example, when fare adjustment module 148 employs a nested approach, the seat protect is assigned to D class (which is the class having a value equal to the EMSR value of \$180) rather than to B class, which has a value of \$900. The nested approach ensures that any booking having a value higher than the probabilistic value of the protected seat will be accepted. However, the nested approach does not take dilution into account; the forecasted class B customer might purchase in class D instead (i.e., the customer might buy down).

[0134] In order to address the potential downsides and/or shortcomings illustrated by the foregoing examples, in various embodiments fare adjustment module 148 is configured to account for, estimate, and/or otherwise model or respond to dilution.

[0135] In a further example shown in Table 4B, operation of fare adjustment module 148 is illustrated.

TABLE 4B

Fare Class	Average Fare	Demand	Cumulative/ Top-Down Demand	Loss if Seat Purchased by higher class	Probability seat purchased by current class	Probability seat purchased by higher class	Expected value accounting for dilution
A	\$300	5	5	N/A	100%	0%	\$300
B	\$200	15	20	-\$100	75% (15/20)	25%	\$125
C	\$100	30	50	-\$100	60% (30/50)	40%	\$ 20
D	\$ 40	50	100	-\$60-	50% (50/100)	50%	-\$ 10

information for each seat in a fare class, in order to determine whether to open a particular fare class, assign seat protects to a particular fare class, and/or the like.

[0129] In various embodiments, fare adjustment module 148 is configured to assign seat protects via a “partitioned” approach. In a partitioned approach, seat protects are assigned based on the class of the EMSR. In other exemplary embodiments, fare adjustment module 148 is configured to assign seat protects via a “nested” approach. In a nested approach, seat protects are assigned based on the value of the EMSR.

[0130] A partitioned approach may calculate the EMSR for a seat based at least in part on the demand forecast and the average fare for each class. A nested approach may calculate the EMSR based on the demand forecast and the average fare of a given class, or based on the joint demand forecast and the joint average fare of a given class and all higher classes.

[0136] In the example scenario illustrated in Table 4B, fare adjustment module 148 determines that having class D open is not a revenue positive decision. Stated differently, the true value of a class D booking (the marginal fare) is the expected value received from a class D booking when class D is open: $\text{marginal fare} = \$40 * (50/100) + (-\$60) * (50/100) = \$-10$. As can be seen, a class D booking results in: $\$40 * (50/100) = \20 of “stimulus”, and $(-\$60) * (50/100) = -\30 in “dilution”.

[0137] In various embodiments, in connection with operation of fare adjustment module 148 and/or forecasting system 115, a seat allocation system is utilized to allocate seats in a class until the joint EMSR value falls below the average fare of the child class. Fare adjustment module 148 may be utilized to modify the seat allocation process for a particular class to account for dilution if a child class is open. After utilization of fare adjustment module 148, a seat allocation system may allocate seats in a class until the joint EMSR

value falls below the adjusted average fare of the child class. With reference now to FIG. 4C, it can be seen that, responsive to operation of fare adjustment module 148, a seat allocation system may assign more seats to a particular class than would otherwise be assigned; stated differently, a seat allocation system may assign more seats to the current class and fewer seats to a child class.

[0138] For example, in a particular example instance illustrated in FIG. 4C, absent operation of fare adjustment module 148, only six seats would be allocated to the current fare class (shown at point A in FIG. 4C). In contrast, responsive to operation of fare adjustment module 148, nine seats are allocated to the current fare class (shown at point B in FIG. 4C). In other words, three additional seats that would otherwise have been allocated to a child class are instead allocated to the current class. In this manner, because additional seats are allocated to the current class and fewer seats are thus available for allocation to a child class, the risk of future buy-down is reduced. In sum, fare adjustment module 148 enables proper valuation of a child class as compared to the parent class, accounting for both stimulation and dilution.

[0139] It will be appreciated that if all passengers were product oriented, prior approaches which presumed independence of demand would be accurate; however, because all passengers are not product oriented, fare adjustment module 148 may be utilized to better quantify the revenue effects of opening a particular fare class for bookings. Moreover, it will be appreciated that fare adjustment module 148 may not necessarily be utilized to set availability or fares for a particular fare class. Rather, fare adjustment module 148 may typically be utilized to determine the true value of the child class for purposes of comparison to the parent class. Revised and/or updated solutions may be generated based at least in part on the output of fare adjustment module 148. Such updated solutions may be utilized by other systems, for example a revenue management system, to set and/or revise actual fare class seat prices for public sale. Moreover, in various exemplary embodiments, the output of fare adjustment module 148 may be utilized as an input to other electronic systems (for example, PROS), for example in order to refine and/or revise a demand forecast and/or optimization.

[0140] In addition to fare adjustment principles, unobscuring principles, and/or unconstraining principles as disclosed hereinabove, forecasting system 115 may be configured to utilize prime class remapping (PCR) module 149. In various exemplary embodiments, PCR module 149 is configured to allocate prime class bookings to fare classes in a manner that results in positive revenue consequences (or reduces and/or minimizes negative revenue consequences).

[0141] Typically, airlines have at least one fare class (referred to herein as a “prime class”) that is given priority over all (or some) other fare classes. Certain prime class contracts also require last seat availability. Because prime class bookings often occur close to a flight departure date, prime class bookings often displace high-revenue fare class bookings. However, prime class bookings are often associated with a revenue amount that is inconsistent with the position they occupy in the class hierarchy. Stated another way, prime class bookings often result in zero additional revenue (for example, prime class bookings arising from use of mileage rewards programs, such as high-priority double mileage redemptions) or minimal additional revenue (certain government and/or corporate bookings), but have high priority. In general, because prime class bookings have a lower average fare than

most fare classes, protecting seats for prime class bookings at the expense of other revenue bookings has a negative impact on airline revenue. Prime class bookings have high priority, but highly variable revenue; accordingly, while airlines are often obligated to protect seats for prime class bookings, doing so typically has a negative effect on revenue.

[0142] Adding complexity to this problem, common forecasting tools in the airline industry are not configured to handle prime class forecasting; a typical approach is simply not to forecast prime class demand at all, but rather to assume prime class demand is zero. However, because prime class bookings commonly occupy from about 1% to about 3% of total bookings, ignoring demand for prime class bookings introduces forecasting errors, resulting, depending on the nature of the errors, in unnecessary buy-down and/or empty (spoiled) seats on a flight.

[0143] Turning now to FIG. 5, in accordance with various embodiments PCR module 149 is configured to implement a method 500 for prime class remapping. Stated generally, PCR module 149 is configured to provide an improved answer to the question: “What is the value of a particular prime class booking?”

[0144] In various embodiments, PCR module 149 is configured to treat a prime class booking as if it were a booking of an equivalent revenue class or higher class. Stated differently, PCR module 149 is configured for “remapping” of prime class bookings into other fare classes, for example for purposes of demand forecasting, seat protection, and/or the like.

[0145] In various embodiments, prime class remapping method 500 utilizes an average fare for prime class bookings. In other exemplary embodiments, method 500 utilizes a known fare for each prime class booking. In certain exemplary embodiments, method 500 employs a hybrid approach, wherein certain prime class bookings are considered to represent an average prime class fare, and certain other prime class bookings are considered to represent a known prime class fare (which may be higher or lower than the average).

[0146] PCR module 149 and/or method 500 may be utilized on any suitable time schedule, as desired. For example, PCR module 149 may be operable after collection of data associated with a DCP, on a weekly basis, on a daily basis, on an hourly basis, on a “triggered” basis (for example, each time a trigger event occurs, such as a prime class booking), and/or the like

[0147] In various embodiments, PCR module 149 identifies the net prime class bookings for a particular flight for the period in question (step 510), as follows: net prime class bookings=current prime class bookings–prior prime class bookings. For each net prime class booking, PCR module 149 identifies the fare class that is closest in value to the value (i.e., average or known) of that net prime class booking (step 520). This closest fare class may be referred to as the “floor class”. However, for the flight in question, at the time of that prime class booking the floor class and/or higher fare classes may have been (or are) closed. If the floor class was or is closed, PCR module 149 identifies an alternative class, the “displacement class” (step 530). The displacement class is the lowest class that is both (i) higher than the floor class, and (ii) currently open for bookings. PCR module 149 “remaps” each prime class booking to its respective target class, which is the higher of the floor class and the displacement class for that prime class booking (step 540). PCR module 149 may also add the remapped net prime bookings to other booking infor-

mation for the current period and/or prior period(s), for example in order to generate cumulative booking information for the current period (step 550). Additionally, if a previously booked prime class booking has been cancelled during the current period, PCR module 149 may “unmap” that prime class booking from the lowest fare class that had previously received any remapped prime class booking (step 560).

[0148] It will be appreciated that any suitable data inputs and/or information may be utilized by PCR module 149, for example historical booking information, estimated fare information for a prime class booking, actual fare information for a prime class booking, currently open fare classes for a flight, prior output of PCR module 149 for a prior time period, departure date information, market information, fare information, and/or the like.

[0149] Operation of PCR module 149 and/or method 500 may be illustrated in the example shown in Table 5A. For a particular example, for Flight 544 during a particular time period, 5 new prime class bookings are received. For 3 of the new prime class bookings (prime class bookings P1, P2, and P3), class K is the floor class; for the remaining 2 prime class bookings (P4 and P5), class L is the floor class. PCR module 149 remaps prime class bookings P1 through P3 into class T (the lowest open class higher than the floor class, class K, which is closed). PCR module 149 remaps prime class bookings P4 and P5 into class L. PCR module 149 outputs the updated remapped prime class booking information in the form of an updated bookings table. The updated bookings table may be used, for example, by other components of forecasting system 115, external system(s) 160, and/or the like.

TABLE 5A

Fare Class	New Prime Class Bookings	Remapped Prime Class Bookings
Prime	5	N/A
Q	N/A	0
N	N/A	0
V	N/A	0
W	N/A	0
L	N/A	2
S	N/A	0
T	N/A	3
G (closed)	N/A	N/A
K (closed)	N/A	N/A
U (closed)	N/A	N/A
E (closed)	N/A	N/A
R (closed)	N/A	N/A

[0150] Principles of operation of PCR module 149 and/or method 500 are further illustrated in the additional example represented in Table 5B. For a particular example Flight 544, during a particular time period, 3 prime class bookings are cancelled, i.e., “new” prime class bookings has the value -3. In this example, 10 prior prime class bookings had previously been remapped into fare classes T, S, L, W, and V, with 2 prime class bookings previously remapped into each class, respectively. -3 prime class bookings were received in the current period, so 2 previously remapped prime class bookings are “unmapped” from fare class T (the lowest fare class that had previously received prime class remapping). Fare class T now holds 0 remapped prime class bookings, so the 1 remaining prime class cancellation is allocated to fare class S (the next lowest fare class that had previously received prime

class remapping). PCR module 149 outputs the updated remapped prime class booking information, for example for use by other components of forecasting system 115, external system(s) 160, and/or the like.

TABLE 5B

Fare Class	“New” Prime Class Bookings	Previously Remapped Prime Class Bookings	Updated Remapped Prime Class Bookings
Prime	-3	N/A	N/A
Q	N/A	0	0
N	N/A	0	0
V	N/A	2	2
W	N/A	2	2
L	N/A	2	2
S	N/A	2	1
T	N/A	2	0
G (closed)	N/A	N/A	N/A
K (closed)	N/A	N/A	N/A
U (closed)	N/A	N/A	N/A
E (closed)	N/A	N/A	N/A
R (closed)	N/A	N/A	N/A

[0151] Via use of PCR module 149 and/or method 500, forecasting system 115 can better account for, modify, and/or forecast prime class demand. Additionally, because prime class bookings have been “remapped” to other fare classes that are suitable for use in external forecasting tools (for example, PROS), forecasting system 115 output may be utilized to incorporate prime class demand into an overall demand model for all fare classes for a particular flight. Via operation of PCR module 149 and/or method 500, forecasting system 115 is also enabled to reduce and/or prevent displacement of high value, late bookings by prime class bookings (which cannot be prevented from booking). Yet further, forecasting system 115 is enabled to more accurately account for the value of a particular prime class booking to an airline. In this manner, forecasting system 115 increases revenues (and/or decreases losses) associated with prime class bookings.

[0152] It will be appreciated that principles of the present disclosure associated with PCR module 149 and/or method 500 are not limited to prime class bookings. More generally, remapping approaches similar to those disclosed herein may suitably be employed to remap (and/or otherwise adjust and/or compensate for) any booking that is capable of being assigned to a fare class having a value that is higher than the value (actual or estimated) of that booking.

[0153] While the exemplary embodiments described herein are described in sufficient detail to enable those skilled in the art to practice principles of the present disclosure, it should be understood that other embodiments may be realized and that logical and/or functional changes may be made without departing from the spirit and scope of the present disclosure. Thus, the detailed description herein is presented for purposes of illustration and not of limitation.

[0154] For the sake of brevity, conventional data management, computer networking, statistical assessment, software application development, and other aspects of exemplary systems and methods (and components thereof) may not be described in detail herein. Furthermore, the connecting lines shown in the various figures contained herein are intended to represent functional relationships and/or physical or communicative couplings between the various elements. It should be

noted that many alternative or additional functional relationships or physical connections may be present in a practical forecasting system.

[0155] While the description references specific technologies, system architectures and data management techniques, practitioners will appreciate that this description is of various embodiments, and that other devices and/or methods may be implemented without departing from the scope of principles of the present disclosure. Similarly, while the description references a user interfacing with the system via a computer user interface, practitioners will appreciate that other interfaces may include mobile devices, kiosks and handheld devices such as mobile phones, smart phones, tablet computing devices, etc.

[0156] While the steps outlined herein represent exemplary embodiments of principles of the present disclosure, practitioners will appreciate that there are any number of computing algorithms and user interfaces that may be applied to create similar results. The steps are presented for the sake of explanation only and are not intended to limit the scope of the present disclosure in any way. Benefits, other advantages, and solutions to problems have been described herein with regard to specific embodiments. However, the benefits, advantages, solutions to problems, and any element(s) that may cause any benefit, advantage, or solution to occur or become more pronounced are not to be construed as critical, required, or essential features or elements of any or all of the claims.

[0157] Systems, methods and computer program products are provided. In the detailed description herein, references to “various embodiments”, “one embodiment”, “an embodiment”, “an example embodiment”, etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to effect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described. After reading the description, it will be apparent to one skilled in the relevant art(s) how to implement principles of the disclosure in alternative embodiments.

[0158] It should be understood that the detailed description and specific examples, indicating exemplary embodiments, are given for purposes of illustration only and not as limitations. Many changes and modifications may be made without departing from the spirit thereof, and principles of the present disclosure include all such modifications. Corresponding structures, materials, acts, and equivalents of all elements are intended to include any structure, material, or acts for performing the functions in combination with other elements. Reference to an element in the singular is not intended to mean “one and only one” unless explicitly so stated, but rather “one or more.” Moreover, when a phrase similar to “at least one of A, B, or C” or “at least one of A, B, and C” is used in the claims or the specification, the phrase is intended to mean any of the following: (1) at least one of A; (2) at least one of B; (3) at least one of C; (4) at least one of A and at least one of B; (5) at least one of B and at least one of C; (6) at least one of A and at least one of C; or (7) at least one of A, at least one of B, and at least one of C.

What is claimed is:

1. A method, comprising:
 - accessing, by a processor for adjusting fares, a bookings table containing seat bookings and forecasted demand for a flight, wherein the seat bookings are grouped into a plurality of fare classes, and wherein at least one fare class in the plurality of fare classes is closed and contains no bookings;
 - determining, by the processor and using the seat bookings and the forecasted demand, the monetary stimulation value of opening the closed class;
 - determining, by the processor and using the seat bookings and the forecasted demand, the monetary dilution penalty of opening the closed class; and
 - comparing, by the processor, the monetary stimulation value to the monetary dilution penalty to generate booking instructions.
2. The method of claim 1, wherein, in response to the monetary stimulation value exceeding the monetary dilution penalty, the booking instructions comprise instructions to open the closed class for bookings.
3. The method of claim 1, wherein, in response to the monetary dilution penalty exceeding the monetary stimulation value, the booking instructions comprise instructions to keep the closed class closed for bookings.
4. The method of claim 1, further comprising communicating, to a reservation system, the booking instructions.
5. The method of claim 1, wherein the determining the monetary stimulation value comprises:
 - determining, by the processor, the proportion of the seat bookings that are product oriented; and
 - multiplying, by the processor, the proportion of the seat bookings that are product oriented by the average fare in the parent class of the closed class.
6. The method of claim 1, wherein the determining the monetary dilution penalty comprises:
 - determining, by the processor, the proportion of the seat bookings that are price oriented; and
 - multiplying, by the processor, the proportion of the seat bookings that are price oriented by a loss associated with a seat booking in the closed class.
7. The method of claim 6, wherein the loss associated with a seat booking in the closed class represents the difference in average fare between the closed class and the parent class.
8. The method of claim 6, wherein the determining the proportion of the seat bookings that are price oriented comprises:
 - identifying, by the processor, any seat bookings in the bookings table that have been modified as a result of an unobscuring algorithm; and
 - identifying, by the processor, any seat bookings in the bookings table that are associated with a price oriented customer.
9. The method of claim 1, further comprising allocating seat protects in a parent fare class until the expected marginal seat revenue (EMSR) in the parent fare class falls below the adjusted average fare of the child fare class.
10. The method of claim 1, further comprising
 - determining, by the processor, which fare classes in the plurality of fare classes have demand that is obscured;
 - calculating, by the processor, the unobscured demand for at least one fare class in the plurality of fare classes;
 - converting, by the processor, the unobscured demand for the at least one fare class into integer values representing seat bookings in the respective fare classes; and

updating, by the processor, the bookings table with the integer values to form an unobscured bookings table.

11. The method of claim 1, further comprising:

accessing, by the processor, the bookings table, wherein the bookings table further contains information representing a net prime class seat booking for a flight, and wherein the net prime class seat booking occurred during a first time period;

determining, by the processor, a floor class for the net prime class seat booking in the bookings table;

determining, by the processor, a displacement class for the net prime class seat booking in the bookings table;

determining, by the processor, a target class for the net prime class seat booking in the bookings table, wherein the target class is the higher of the floor class and the displacement class; and

remapping, by the processor, the net prime class seat booking in the bookings table into its target class to form a remapped bookings table.

12. A non-transitory computer-readable storage medium having computer-executable instructions stored thereon that, in response to execution by a processor for adjusting fares, causes the processor to perform operations comprising:

accessing, by the processor, a bookings table containing seat bookings and forecasted demand for a flight, wherein the seat bookings are grouped into a plurality of fare classes, and wherein at least one fare class in the plurality of fare classes is closed and contains no bookings;

determining, by the processor and using the seat bookings and the forecasted demand, the monetary stimulation value of opening the closed class;

determining, by the processor and using the seat bookings and the forecasted demand, the monetary dilution penalty of opening the closed class; and

comparing, by the processor, the monetary stimulation value to the monetary dilution penalty to generate booking instructions.

* * * * *