

FORECASTING DATA WITH REAL-TIME UPDATES

CROSS-REFERENCE TO RELATED APPLICATIONS

This application claims priority to U.S. Provisional Patent Application serial number 60/565,758, filed April 26, 2004.

BACKGROUND

[0001] Data is used by the investment and business communities to judge the value of investments, the business case for transactions, the performance of managers, and trends in industries, among many other things. Data may be used in other areas and by other communities to make judgments and decisions on a variety of matters. As such, data for a company or organization is important in general, and up-to-date data (such as projected sales data for example) is potentially invaluable. However, forecast data in a company is typically stored in formats or systems which are not amenable to updates on an asynchronous or random basis.

[0002] For example, forecasting cycles are often based on monthly and quarterly updates to information. Government regulations often require reporting on no more than a quarterly basis (every three months). Activities in the real world rarely occur on scheduled dates for updates, a customer may cancel or enhance an order at any time. Moreover, indirect actions with direct effects on customers (such as competitor product announcements or vendor supply changes for example) are also rarely coordinated with a time which is convenient based on accounting schedules.

[0003] Thus, it may be advantageous to provide a system which allows for updates on a random or asynchronous basis. Additionally, information in the form of projections is often based on judgment. Thus, it may be advantageous to provide a system in which changes may be made to information based on judgments made after input of such information. Moreover, some information (datapoints) may have particular significance. Thus, it may be advantageous to provide a system in which a user may be notified of changes to particular datapoints.

SUMMARY

[0004] The present invention is described and illustrated in conjunction with systems, apparatuses and methods of varying scope. In addition to the aspects of the present invention described in this summary, further aspects of the invention will become apparent by reference to the drawings and by reading the detailed description that follows. A method and apparatus for forecasting data with real-time updates is described.

[0005] In one embodiment, the invention is a system. The system includes a first client. The system also includes an analysis server coupled to the first client. The system further includes a first customer database of information coupled to the analysis server. The first customer database is to embody forecast data and to receive essentially real-time updates to the forecast data. The first customer database may support an OLAP cube associated with the analysis server.

[0006] The analysis server may incorporate an OLAP cube therein, with the OLAP cube to analyze and update the information of the first customer database. The first customer database may be dedicated to use by a first set of selected users of the system. An identification server coupled to the first client and coupled to the first customer database may also be included. The first client may be a smart client, and the first client may include its own OLAP cube. A second client may also be coupled to the identification server and the analysis server. Moreover, a second customer database of information may be coupled to the identification server and the analysis server, with the second customer database dedicated to use by a second set of users of the system. The information manipulated by the system may be financial information.

[0007] The first customer database may include watches of data. The first customer database may also include sales forecast data. The first customer database may further include sales forecast data with changes from non-sales personnel.

[0008] In another embodiment, the invention is also a system. The system includes an analysis server including an OLAP cube. The system also includes an information database coupled to the analysis server to support the OLAP cube. The information database is to embody forecast data and to receive essentially real-time updates to the forecast data. The system also may include an identification server coupled to the information database. The system may further include a client coupled to the analysis server and to the identification server. Moreover, the information database may store financial information.

[0009] In yet another embodiment, the invention is a method of maintaining information. The method includes receiving a set of forecast data. The method also includes incorporating the forecast data into a database of the information through an OLAP cube. The method further includes extracting a baseline forecast from the database. Also, the method includes receiving updates to the database. Moreover, the method includes propagating updates almost immediately throughout the information through the OLAP cube.

[0010] The method may also include watching a set of watched data points of the information. The method may further include notifying a user of changes in the set of watched data points responsive to the propagating and receiving updates.

Similarly, the method may also include providing information to a user. The method may further include receiving expected changes of the information from the user. The method may also include propagating the expected changes as updates almost immediately throughout the information through the OLAP cube.

[0011] Additionally, the method may include reviewing updates received and propagated through the information. The method may include providing user-readable updates of information responsive to the reviewing. The forecast data may be financial information. Also, the method may be executed by a processor in response to instructions, with the instructions embodied in a machine-readable medium. Moreover, the method may include receiving actual data corresponding to the information. The method may further include comparing the actual data to the information and providing user-readable comparisons of information and the actual data.

[0012] In one embodiment, the invention is a method. The method includes receiving forecast information from sales people in a computer. The method further includes receiving comments on specific entries of the forecast information from non-sales people in the computer. The method also includes receiving changes of the forecast information from the non-sales people in the computer. The method additionally includes providing a display of the comments and the changes to the sales people in an interface to the computer.

[0013] The comments may be received from marketing and executive personnel. Forecast information and associated comments and changes may be stored in a database and OLAP cube associated with the computer. The computer may actually be

a network of computers hosting the database and OLAP cube. The method may be executed by a processor in response to instructions, with the instructions embodied in a machine-readable medium.

[0014] The method may further include receiving requests to watch specific data of the forecast information. Moreover, the method may include providing updates when specific data changes responsive to the requests to watch. Similarly, the method may include receiving judgments from the non-sales people related to a set of data of the forecast information and changing the set of data of the forecast responsive to the judgments. The method may also include analyzing the changes for largest magnitude changes. The method may further include providing a set of impacts to the sales people, the impacts showing the changes based on a rank ordering of magnitude of the changes.

[0015] Moreover, the comments and changes may be received through devices other than the computer that are coupled to the computer for communication. Similarly, the forecast data may be received through devices other than the computer that are coupled to the computer for communication. Additionally, the interface to the computer may be through devices other than the computer that are coupled to the computer for communication.

[0016] In yet another embodiment, the invention is a system. The system includes a user interface to receive forecast data. The system also includes an analysis server including an OLAP cube. The system further includes an information database coupled to the analysis server to support the OLAP cube. The information database is to

embody the forecast data and to receive essentially real-time updates to the forecast data through the user interface. The user interface is to display the forecast data through the user interface with feedback as the data changes over time.

[0017] In still another embodiment, the invention is a method of providing feedback from members of an organization on forecasts to sales people entering data for the forecasts. The method includes receiving the data for the forecasts from the sales people in a computer. The method also includes storing the data for the forecasts within a database in conjunction with an OLAP cube as forecast information. The method further includes receiving comments on specific entries of the forecast information in the OLAP cube from the members of the organization in the computer. The method also includes receiving changes of the forecast information from the members of the organization in the computer. Moreover, the method includes integrating the comments and the changes into the OLAP cube and the database. Also, the method includes providing a display of the comments and the changes to the sales people in an interface to the computer.

[0018] In still another embodiment, the invention is a method. The method includes receiving forecast data in a database with an associated window value. The method also includes accessing data based on associated window values. The method further includes comparing the data accessed based on associated window values to other data.

[0019] In another embodiment, the invention is a method of maintaining information. The method includes receiving a set of forecast data. The method also

includes incorporating the forecast data into a database of the information through an OLAP cube. The method further includes extracting a baseline forecast from the database. The method also includes receiving updates to the database. The method further includes propagating updates almost immediately throughout the information through the OLAP cube. The method also includes extracting an updated baseline forecast from the database, the updated baseline forecast derived from the baseline forecast.

[0020] In another embodiment, the invention is a method. The method includes receiving forecast information in a computer from a first group of users. The method also includes integrating the forecast information into a database and corresponding OLAP cube accessible by the computer. The method further includes receiving comments on specific entries of the forecast information in the computer from a second group of users. The method also includes receiving changes of the forecast information in the computer from a second group of users. The method further includes integrating comments and changes into the database and OLAP cube. Moreover, the method includes providing a display of the comments and the changes to the first group of users in an interface to the computer.

[0021] The present invention is exemplified in the various embodiments described, and is limited in spirit and scope only by the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0022] The present invention is illustrated in various exemplary embodiments and is limited in spirit and scope only by the appended claims.

[0023] Fig. 1 illustrates information flowing through an organization.

[0024] Fig. 2 illustrates information for a sales person in one embodiment.

[0025] Fig. 3 illustrates information for a customer in one embodiment.

[0026] Fig. 4 illustrates information for a region in one embodiment.

[0027] Fig. 5 illustrates information for a larger region in one embodiment.

[0028] Fig. 6 illustrates information for worldwide operations of a company in one embodiment.

[0029] Fig. 7 illustrates an embodiment of a network which may be used in conjunction with forecasting data with real-time updates.

[0030] Fig. 8 illustrates an embodiment of a machine which may be used in conjunction with forecasting data with real-time updates.

[0031] Fig. 9 illustrates an embodiment of a system for use in forecasting data with real-time updates.

[0032] Fig. 10 illustrates an alternate embodiment of a system for use in forecasting data with real-time updates.

[0033] Fig. 11 illustrates an embodiment of a method for use in forecasting data with real-time updates.

[0034] Fig. 12 illustrates an alternate embodiment of a method for use in forecasting data with real-time updates.

[0035] Fig. 13 illustrates display by product of information in an embodiment.

[0036] Fig. 14 illustrates display by customer of information in an embodiment.

[0037] Fig. 15 illustrates display of a specific impact in an embodiment.

[0038] Fig. 16 illustrates an operations report for information in an embodiment.

[0039] Fig. 17 illustrates an inventory report for information in an embodiment.

[0040] Fig. 18 illustrates a gap report for information in an embodiment.

[0041] Fig. 19 illustrates a drilled down display of information for a sales person in one embodiment.

[0042] Fig. 20 illustrates an updated forecast in an embodiment.

[0043] Fig. 21 illustrates a drilled down display of a region in an embodiment.

[0044] Fig. 22 illustrates application of judgment in an embodiment.

[0045] Fig. 23 further illustrates application of judgment in an embodiment.

[0046] Fig. 24 illustrates a forecast after application of judgment in an embodiment.

[0047] Fig. 25 illustrates change history for information in an embodiment.

[0048] Fig. 26 illustrates addition of a watch in an embodiment.

[0049] Fig. 27 illustrates an embodiment of a process of providing feedback in a real-time forecasting system.

[0050] Fig. 28 illustrates relationships between various groups and data associated with those groups in an embodiment.

[0051] Fig. 29 illustrates another embodiment of a process of providing feedback in a real-time forecasting system.

[0052] Fig. 30 illustrates a real-time forecasting system and its relationship with various groups in an embodiment.

[0053] Fig. 31 illustrates an embodiment of a process of providing sales feedback in a real-time forecasting system.

[0054] Fig. 32 illustrates changes to an atomic data point in one embodiment.

[0055] Fig. 33 illustrates an embodiment of a process of generating a baseline.

[0056] Fig. 34 illustrates an embodiment of a process of operating a real-time forecasting system in conjunction with a baseline.

[0057] Fig. 35 illustrates relationships between baselines and potential operations in an embodiment.

[0058] Fig. 36 illustrates an embodiment of a process of gathering data in a real-time forecasting system.

[0059] Fig. 37 illustrates an embodiment of a process of annotating data in a real-time forecasting system.

[0060] Fig. 38 illustrates an embodiment of a process of displaying annotated data in a real-time forecasting system.

[0061] Fig. 39 illustrates an embodiment of a process of revising annotated data in a real-time forecasting system.

[0062] Fig. 40 illustrates an embodiment of a process of operating a real-time forecasting system.

[0063] Fig. 41 illustrates a timeline of operations in conjunction with a real-time forecasting system in one embodiment.

[0064] Fig. 42 illustrates an embodiment of a data structure.

[0065] Fig. 43 illustrates an embodiment of an OLAP structure.

[0066] Fig. 44 illustrates an embodiment of a process of making changes in an OLAP structure.

[0067] Fig. 45 illustrates an embodiment of a process of responding to a query in an OLAP structure.

[0068] Fig. 46 illustrates an embodiment of a data structure in conjunction with a timeline.

[0069] Fig. 47 illustrates an embodiment of a process of responding to a request for data.

[0070] Fig. 48 illustrates an embodiment of a process of storing a change in data.

[0071] Fig. 49 illustrates an embodiment of a set of baselines and tracked data.

[0072] Fig. 50 illustrates an embodiment of a medium and associated devices.

[0073] Like reference symbols in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0074] The present invention is described and illustrated in conjunction with systems, apparatuses and methods of varying scope. In addition to the aspects of the present invention described in this summary, further aspects of the invention will become apparent by reference to the drawings and by reading the detailed description that follows. A method and apparatus for forecasting data with real-time updates is described. In general, the method and apparatus relate to gathering forecast data from a variety of sources, developing a baseline forecast from the gathered data, and updating the baseline forecast based on essentially real-time changes in data as gathered from the variety of sources and other data sources. Moreover, the method and apparatus allow for viewing of forecast data with updates and may allow for simulation or alteration of the data.

[0075] In one embodiment, the invention is a system. The system includes a first client. The system also includes an analysis server coupled to the first client. The system further includes a first customer database of information coupled to the analysis server. The first customer database is to embody forecast data and to receive essentially real-time updates to the forecast data. The first customer database may support an OLAP cube associated with the analysis server.

[0076] The analysis server may incorporate an OLAP cube therein, with the OLAP cube to analyze and update the information of the first customer database. The first customer database may be dedicated to use by a first set of selected users of the system. An identification server coupled to the first client and coupled to the first customer database may also be included. The first client may be a smart client, and the

first client may include its own OLAP cube. A second client may also be coupled to the identification server and the analysis server. Moreover, a second customer database of information may be coupled to the identification server and the analysis server, with the second customer database dedicated to use by a second set of users of the system. The information manipulated by the system may be financial information.

[0077] The first customer database may include watches of data. The first customer database may also include sales forecast data. The first customer database may further include sales forecast data with changes from non-sales personnel.

[0078] In another embodiment, the invention is also a system. The system includes an analysis server including an OLAP cube. The system also includes an information database coupled to the analysis server to support the OLAP cube. The information database is to embody forecast data and to receive essentially real-time updates to the forecast data. The system also may include an identification server coupled to the information database. The system may further include a client coupled to the analysis server and to the identification server. Moreover, the information database may store financial information.

[0079] In yet another embodiment, the invention is a method of maintaining information. The method includes receiving a set of forecast data. The method also includes incorporating the forecast data into a database of the information through an OLAP cube. The method further includes extracting a baseline forecast from the database. Also, the method includes receiving updates to the database. Moreover, the method includes propagating updates almost immediately throughout the information

through the OLAP cube.

[0080] The method may also include watching a set of watched data points of the information. The method may further include notifying a user of changes in the set of watched data points responsive to the propagating and receiving updates. Similarly, the method may also include providing information to a user. The method may further include receiving expected changes of the information from the user. The method may also include propagating the expected changes as updates almost immediately throughout the information through the OLAP cube.

[0081] Additionally, the method may include reviewing updates received and propagated through the information. The method may include providing user-readable updates of information responsive to the reviewing. The forecast data may be financial information. Also, the method may be executed by a processor in response to instructions, with the instructions embodied in a machine-readable medium. Moreover, the method may include receiving actual data corresponding to the information. The method may further include comparing the actual data to the information and providing user-readable comparisons of information and the actual data.

[0082] In one embodiment, the invention is a method. The method includes receiving forecast information from sales people in a computer. The method further includes receiving comments on specific entries of the forecast information from non-sales people in the computer. The method also includes receiving changes of the forecast information from the non-sales people in the computer. The method additionally

includes providing a display of the comments and the changes to the sales people in an interface to the computer.

[0083] The comments may be received from marketing and executive personnel. Forecast information and associated comments and changes may be stored in a database and OLAP cube associated with the computer. The computer may actually be a network of computers hosting the database and OLAP cube. The method may be executed by a processor in response to instructions, with the instructions embodied in a machine-readable medium.

[0084] The method may further include receiving requests to watch specific data of the forecast information. Moreover, the method may include providing updates when specific data changes responsive to the requests to watch. Similarly, the method may include receiving judgments from the non-sales people related to a set of data of the forecast information and changing the set of data of the forecast responsive to the judgments. The method may also include analyzing the changes for largest magnitude changes. The method may further include providing a set of impacts to the sales people, the impacts showing the changes based on a rank ordering of magnitude of the changes.

[0085] Moreover, the comments and changes may be received through devices other than the computer that are coupled to the computer for communication. Similarly, the forecast data may be received through devices other than the computer that are coupled to the computer for communication. Additionally, the interface to the

computer may be through devices other than the computer that are coupled to the computer for communication.

[0086] In yet another embodiment, the invention is a system. The system includes a user interface to receive forecast data. The system also includes an analysis server including an OLAP cube. The system further includes an information database coupled to the analysis server to support the OLAP cube. The information database is to embody the forecast data and to receive essentially real-time updates to the forecast data through the user interface. The user interface is to display the forecast data through the user interface with feedback as the data changes over time.

[0087] In still another embodiment, the invention is a method of providing feedback from members of an organization on forecasts to sales people entering data for the forecasts. The method includes receiving the data for the forecasts from the sales people in a computer. The method also includes storing the data for the forecasts within a database in conjunction with an OLAP cube as forecast information. The method further includes receiving comments on specific entries of the forecast information in the OLAP cube from the members of the organization in the computer. The method also includes receiving changes of the forecast information from the members of the organization in the computer. Moreover, the method includes integrating the comments and the changes into the OLAP cube and the database. Also, the method includes providing a display of the comments and the changes to the sales people in an interface to the computer.

[0088] In still another embodiment, the invention is a method. The method includes receiving forecast data in a database with an associated window value. The method also includes accessing data based on associated window values. The method further includes comparing the data accessed based on associated window values to other data.

[0089] In another embodiment, the invention is a method of maintaining information. The method includes receiving a set of forecast data. The method also includes incorporating the forecast data into a database of the information through an OLAP cube. The method further includes extracting a baseline forecast from the database. The method also includes receiving updates to the database. The method further includes propagating updates almost immediately throughout the information through the OLAP cube. The method also includes extracting an updated baseline forecast from the database, the updated baseline forecast derived from the baseline forecast.

[0090] In another embodiment, the invention is a method. The method includes receiving forecast information in a computer from a first group of users. The method also includes integrating the forecast information into a database and corresponding OLAP cube accessible by the computer. The method further includes receiving comments on specific entries of the forecast information in the computer from a second group of users. The method also includes receiving changes of the forecast information in the computer from a second group of users. The method further includes integrating comments and changes into the database and OLAP cube. Moreover, the

method includes providing a display of the comments and the changes to the first group of users in an interface to the computer.

[0091] Fig. 1 illustrates information flowing through an organization. Organization 100 may be typical of companies or other organizations concerned with finances. Ceo 110 needs financial information as a basic part of the job. Sales organization 150 is thus asked to prepare a forecast of sales. This forecast may span months or years for example. Examples illustrated are for 6 months, but other time periods may fit circumstances in various situations. Other forecasts may be requested within an organization within the spirit and scope of the present invention. For example, forecasts of inventory or expenses may be prepared and tracked in real-time.

[0092] Within sales organization 150, sales representatives 175 provide forecasts of their upcoming sales. Similarly, independent representative 180 and distributor 190 provide forecasts of upcoming sales. Area manager 170 receives these forecasts, and passes them up to vice president of sales 160. Vp 160 then passes the forecasts to marketing department 140. At this point, and at previous links, feedback or analysis of the financial data (forecasts) may occur, such as through changes to estimates, requests for information about or verification of data, or other forms of feedback or analysis.

[0093] Marketing 140 then sends current data to production 130 (engineering and manufacturing for example). Production 130 may comment and provide changes based on manufacturing considerations (such as delays or stockpiles for example) and then pass the information to finance 120. Finance department 120 may

comment and provide changes based on financial considerations, such as availability of capital or status of accounts (such as past due accounts for example). Finance department 120 then passes the updated forecast data to ceo 110 as a baseline forecast.

[0094] Ceo 110 may use this baseline for managerial analysis and for reference when speaking to non-members of the organization, such as news media outlets, customers, vendors and regulators for example. With a static forecast, the data may be stale by the time ceo 110 sees it. With real-time updates, ceo 110 may rely more effectively on available data to analyze and comment on the organization's financial situation.

[0095] To illustrate in further detail the forecasting and update process, reference may be made to how data is provided initially. Again, the process is presented in terms of sales data, but data of various types may be forecasted and tracked in real-time. Fig. 2 illustrates information for a sales person in one embodiment. The information is entered and/or displayed through user interface 200. Field or frame 260 is a display of information for the sales person 220. As illustrated, this is a display in currency (such as dollars for example) for a salesperson 220 including a company 230, a part 240 for that company, details 245 for the part 240, and another part 250 for which details are not presently selected. Field or frame 260 is a user interface for products sold by salesperson 220 which allows for entry of forecast data related to specific products for the salesperson 220.

[0096] In one embodiment, a sales representative or similar individual (user) enters information into each cell in frame 260, and is required to "touch" each cell

(enter or confirm data in the cell) to attempt to verify that no data is inadvertently left out or entered incorrectly. Moreover, the user may be required to touch each cell of the summary data of part 270. Additionally, status information related to what is being entered is displayed as status 205, and submit 210 and exit 215 buttons are provided for submission of entered data and exit of the software respectively. Once data has been entered, a similar user interface may be used to display the data. If changes are made to the data, those changes may also be displayed as described below.

[0097] As illustrated, a similar user interface 300 may be used for display of information once it is entered. Fig. 3 illustrates information for a customer in one embodiment. Interface 300 provides a forecast overview, impact messages, top 10 customers, and navigation tools. Forecast display 310 provides information about a particular sales representative (for all sales people) in one embodiment. The information is displayed in a cell format, with sales people separated by row and columns devoted to time periods. Display 360 indicates what is being displayed, in this case a representative sales inc. In the ne usa part of the east usa region. Moreover, selector 365 determines the format of the display 310, in this case a summary, time selector 370 determines the displayed time period, and type selector 375 determines the type of display, such as dollars or units. Additionally, display 310 provides for user selection of a method of viewing (by months as illustrated) and a type of view (table as illustrated).

[0098] Another part of the display is the impacts message display 315, which provides messages about impacts to a projection based on changes. Still another part of the display is top 10 customers display 305, which may be used to provide

forecasts on the top 10 customers in real time, regardless of what else is displayed. Additionally, status and navigation tools are provided. Forecast button 335 leads to the displayed forecast data. Drill down button 330 allows a user to delve into details of an entry of a subset of displayed data. Settings button 325 allows the user to change settings of the display. Help button 320 allows the user to access online help and potentially to access help over a network for example. Identity 345 displays an identity of the current user, and projection status 340 displays the status of the projection (such as whether it needs to be approved or it is active and will provide updates). Moreover, logout button 355 and home button 350 allow for exiting the system or navigating to a predetermined home part of the system respectively.

[0099] With information from users related to various customers and areas, an overview of a broader area may be provided. Fig. 4 illustrates information for a region in one embodiment. By navigating to a different region, or a region encompassing the previously displayed data for example, information for a region may be displayed. In this example, the display of frame 310 now provides data for sales inc., as the data for the ne usa region. Display 360 indicates what region is displayed. Note that the data displayed for sales inc. Is an aggregation of the data displayed in fig. 3, as this is essentially displaying data at a higher level of abstraction or a different level of organization from the data of fig. 3.

[00100] Aggregation of data for larger regions may similarly proceed. Fig. 5 illustrates information for a larger region in one embodiment. Display 310 now provides data for the ne usa and se usa regions, with the ne usa region aggregating the

sales inc. Data of fig. 4. Similarly, the se usa region aggregates appropriate data. As may be expected, display 360 indicates what is being displayed, in this case the east usa area.

[00101] As one may expect, aggregation may ultimately go to a worldwide level. Fig. 6 illustrates information for worldwide operations of a company in one embodiment. User interface 300 now provides data on a worldwide basis, with an indication of what area is displayed in display 360. In particular, display 310 provides data for east usa, west usa, and asia pacific regions. The east usa data is an aggregation of the data of fig. 5. Moreover, as illustrated, no changes have been made to the information. Additionally, the display 310 allows for display by area or entity (tab 380), by customer (tab 385), and by product (tab 390). However, having gathered the information and displayed it, it may be useful to manipulate the information, such as by various users over a network using individual clients or workstations for example.

[00102] The following description of figs. 7-8 is intended to provide an overview of computer hardware and other operating components suitable for performing the methods of the invention described above and hereafter, but is not intended to limit the applicable environments. Similarly, the computer hardware and other operating components may be suitable as part of the apparatuses of the invention described above. The invention can be practiced with other computer system configurations, including hand-held devices, multiprocessor systems, microprocessor-based or programmable consumer electronics, network pcs, minicomputers, mainframe computers, and the like. The invention can also be practiced in distributed computing environments where tasks

are performed by remote processing devices that are linked through a communications network.

[00103] Fig. 7 shows several computer systems that are coupled together through a network 705, such as the internet. The term "internet" as used herein refers to a network of networks which uses certain protocols, such as the tcp/ip protocol, and possibly other protocols such as the hypertext transfer protocol (http) for hypertext markup language (html) documents that make up the world wide web (web). The physical connections of the internet and the protocols and communication procedures of the internet are well known to those of skill in the art.

[00104] Access to the internet 705 is typically provided by internet service providers (isp), such as the isps 710 and 715. Users on client systems, such as client computer systems 730, 740, 750, and 760 obtain access to the internet through the internet service providers, such as isps 710 and 715. Access to the internet allows users of the client computer systems to exchange information, receive and send e-mails, and view documents, such as documents which have been prepared in the html format. These documents are often provided by web servers, such as web server 720 which is considered to be "on" the internet. Often these web servers are provided by the isps, such as isp 710, although a computer system can be set up and connected to the internet without that system also being an isp.

[00105] The web server 720 is typically at least one computer system which operates as a server computer system and is configured to operate with the protocols of the world wide web and is coupled to the internet. Optionally, the web

server 720 can be part of an isp which provides access to the internet for client systems. The web server 720 is shown coupled to the server computer system 725 which itself is coupled to web content 795, which can be considered a form of a media database. While two computer systems 720 and 725 are shown in fig. 7, the web server system 720 and the server computer system 725 can be one computer system having different software components providing the web server functionality and the server functionality provided by the server computer system 725 which will be described further below.

[00106] Client computer systems 730, 740, 750, and 760 can each, with the appropriate web browsing software, view html pages provided by the web server 720. The isp 710 provides internet connectivity to the client computer system 730 through the modem interface 735 which can be considered part of the client computer system 730. The client computer system can be a personal computer system, a network computer, a web tv system, or other such computer system.

[00107] Similarly, the isp 715 provides internet connectivity for client systems 740, 750, and 760, although as shown in fig. 7, the connections are not the same for these three computer systems. Client computer system 740 is coupled through a modem interface 745 while client computer systems 750 and 760 are part of a lan. While fig. 7 shows the interfaces 735 and 745 as generically as a "modem," each of these interfaces can be an analog modem, isdn modem, cable modem, satellite transmission interface (e.g. "direct pc"), or other interfaces for coupling a computer system to other computer systems.

[00108] Client computer systems 750 and 760 are coupled to a lan 770 through network interfaces 755 and 765, which can be ethernet network or other network interfaces. The lan 770 is also coupled to a gateway computer system 775 which can provide firewall and other internet related services for the local area network. This gateway computer system 775 is coupled to the isp 715 to provide internet connectivity to the client computer systems 750 and 760. The gateway computer system 775 can be a conventional server computer system. Also, the web server system 720 can be a conventional server computer system.

[00109] Alternatively, a server computer system 780 can be directly coupled to the lan 770 through a network interface 785 to provide files 790 and other services to the clients 750, 760, without the need to connect to the internet through the gateway system 775.

[00110] Fig. 8 shows one example of a conventional computer system that can be used as a client computer system or a server computer system or as a web server system. Such a computer system can be used to perform many of the functions of an internet service provider, such as isp 710. The computer system 800 interfaces to external systems through the modem or network interface 820. It will be appreciated that the modem or network interface 820 can be considered to be part of the computer system 800. This interface 820 can be an analog modem, isdn modem, cable modem, token ring interface, satellite transmission interface (e.g. "direct pc"), or other interfaces for coupling a computer system to other computer systems.

[00111] The computer system 800 includes a processor 810, which can be a conventional microprocessor such as an intel pentium microprocessor or motorola power pc microprocessor. Memory 840 is coupled to the processor 810 by a bus 870. Memory 840 can be dynamic random access memory (dram) and can also include static ram (sram). The bus 870 couples the processor 810 to the memory 840, also to non-volatile storage 850, to display controller 830, and to the input/output (i/o) controller 860.

[00112] The display controller 830 controls in the conventional manner a display on a display device 835 which can be a cathode ray tube (crt) or liquid crystal display (lcd). The input/output devices 855 can include a keyboard, disk drives, printers, a scanner, and other input and output devices, including a mouse or other pointing device. The display controller 830 and the i/o controller 860 can be implemented with conventional well known technology. A digital image input device 865 can be a digital camera which is coupled to an i/o controller 860 in order to allow images from the digital camera to be input into the computer system 800.

[00113] The non-volatile storage 850 is often a magnetic hard disk, an optical disk, or another form of storage for large amounts of data. Some of this data is often written, by a direct memory access process, into memory 840 during execution of software in the computer system 800. One of skill in the art will immediately recognize that the terms "machine-readable medium" or "computer-readable medium" includes any type of storage device that is accessible by the processor 810 and also encompasses a carrier wave that encodes a data signal.

[00114] The computer system 800 is one example of many possible computer systems which have different architectures. For example, personal computers based on an intel microprocessor often have multiple buses, one of which can be an input/output (i/o) bus for the peripherals and one that directly connects the processor 810 and the memory 840 (often referred to as a memory bus). The buses are connected together through bridge components that perform any necessary translation due to differing bus protocols.

[00115] Network computers are another type of computer system that can be used with the present invention. Network computers do not usually include a hard disk or other mass storage, and the executable programs are loaded from a network connection into the memory 840 for execution by the processor 810. A web tv system, which is known in the art, is also considered to be a computer system according to the present invention, but it may lack some of the features shown in fig. 8, such as certain input or output devices. A typical computer system will usually include at least a processor, memory, and a bus coupling the memory to the processor.

[00116] In addition, the computer system 800 is controlled by operating system software which includes a file management system, such as a disk operating system, which is part of the operating system software. One example of an operating system software with its associated file management system software is the family of operating systems known as windows® from microsoft corporation of redmond, washington, and their associated file management systems. Another example of an operating system software with its associated file management system software is the

linux operating system and its associated file management system. The file management system is typically stored in the non-volatile storage 850 and causes the processor 810 to execute the various acts required by the operating system to input and output data and to store data in memory, including storing files on the non-volatile storage 850.

[00117] Some portions of the detailed description are presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the means used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. It has proven convenient at times, principally for reasons of common usage, to refer to these signals as bits, values, elements, symbols, characters, terms, numbers, or the like.

[00118] It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the following discussion, it is appreciated that throughout the description, discussions utilizing terms such as "processing" or "computing" or "calculating" or "determining" or "displaying" or the like, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data

represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

[00119] The present invention, in some embodiments, also relates to apparatus for performing the operations herein. This apparatus may be specially constructed for the required purposes, or it may comprise a general purpose computer selectively activated or reconfigured by a computer program stored in the computer. Such a computer program may be stored in a computer readable storage medium, such as, but is not limited to, any type of disk including floppy disks, optical disks, cd-roms, and magnetic-optical disks, read-only memories (roms), random access memories (rams), eproms, eeproms, magnetic or optical cards, or any type of media suitable for storing electronic instructions, and each coupled to a computer system bus.

[00120] The algorithms and displays presented herein are not inherently related to any particular computer or other apparatus. Various general purpose systems may be used with programs in accordance with the teachings herein, or it may prove convenient to construct more specialized apparatus to perform the required method steps. The required structure for a variety of these systems will appear from the description below. In addition, the present invention is not described with reference to any particular programming language, and various embodiments may thus be implemented using a variety of programming languages.

[00121] Various networks and machines such as those illustrated in figs. 7 and 8 may be utilized. Fig. 9 illustrates an embodiment of a system for use in forecasting data with real-time updates. System 900 includes clients, servers, and supporting databases. Various embodiments of systems may be used, with different configurations as needed due to circumstances surrounding an implementation or installation of such a system. The system 900 may be used with various types of data which is suitable for both forecasting and for essentially real-time updates. Essentially real-time updates refer to updates provided to the system as soon as practicable – this may be immediate, or it may occur as soon as a person with the appropriate information is in a position to relay that information to the system 900.

[00122] System 900 includes a client portion, with clients 910 and 970, a server portion with identification server 920 and analysis server 940, and a database portion with databases 930, 950 and 960. In one embodiment, databases 930, 950 and 960 are each dedicated to specific customers (such as a first customer, second customer and third customer). Identification or authentication server 920 received access requests from various clients such as clients 910 and 970. Server 920 then authenticates or identifies the client(s) and current users to determine which database (if any) should be accessible. Analysis server 940 then starts receiving requests from the clients, allowing for analysis of data in the selected database.

[00123] Authentication and access may be handled in various ways. For example, once a client (client 910 for example) is authenticated, it may be redirected to the analysis server 940 with something such as a token which encodes an address for

server 940 and information about which database (such as 930) to use. Alternatively, client 910 may receive a key (such as a portion of a public key-private key pair for example), which may be used to access a previously known address for server 940 (the key may enable a response for example). The key may also be required to be transmitted from server 940 to database 930 to access data, for example.

[00124] Analysis server 940 may be implemented in part using an olap cube. Olap cubes are available from various commercial entities, such as microsoft corp. Of redmond, washington, for example. An olap cube may perform automated analysis of data when supported by a database such as database 930 for example, allowing for fast throughput of data and fast propagation of changes. In one embodiment, all analysis occurs at server 940, as client 910 is a local client used only to submit information and queries and to view information in the user interface. In another embodiment, client 970 is a smart client which incorporates some analysis capabilities (such as through a local olap cube and repository for example). Client 970, as a smart client, can then be operated in isolation from the network and server 940, while still allowing for analysis and display of data actually stored or replicated at client 970.

[00125] As illustrated, fig. 9 reflects a system which may be distributed geographically and organizationally, thus allowing for revenue generation based on access to and maintenance of the system for example. Fig. 10 illustrates an alternate embodiment of a system for use in forecasting data with real-time updates. Site 1000 includes a client 1010 (potentially many clients), a server 1020 and a database 1030. As the system is self-contained, authentication may or may not be needed (site 1000 may be

restricted to authorized users for example). Client 1010 may be implemented as a local client without analytical capabilities, or as a smart client, and server 1020 may be implemented for analysis with an olap cube for example.

[00126] How the systems and other embodiments operate may vary. Fig. 11 illustrates an embodiment of a method for use in forecasting data with real-time updates. The method (1100) and other methods of this document are illustrated as a set of process modules which may be rearranged and may be performed or operated in a parallel or series manner, for example. At module 1110, data is received, such as financial data for an initial forecast. At module 1120, the data is propagated in the system (such as in an olap cube or supporting database). At module 1130, a forecast baseline is formed, such as due to arriving at a deadline or due to a determination that enough data has been collected for example.

[00127] With the baseline available, some information about the forecast may change, and other information may remain static. At module 1140, changes are received. At module 1150, the received changes are propagated or processed, with updates propagating through the system, alerts or notifications being sent, and any recorded adjustments being applied as appropriate. Modules 1140 and 1150 may be executed multiple times in an iterative fashion as changes are received, thus allowing for essentially continuous and almost real-time updates. Moreover, as time comes for the next forecast (the next quarter for example), the process may return to module 1110.

[00128] Specific processes may be utilized in some embodiments. Fig. 12 illustrates an alternate embodiment of a method for use in forecasting data with real-time

updates. At module 1210, initial data of a forecast is received, such as raw data from sales representatives for example. At module 1220, the data received is propagated through the system. At module 1230, revisions to data, such as judgments applied to raw data leading up to a forecast are received. Such judgments may come from supervisors of sales representatives, marketing department personnel, manufacturing or engineering personnel, financial personnel, and even the ceo or other high-level personnel. Modules 1210, 1220 and 1230 may be repeated due to data arriving at different times or other variations in input cycles for example. At module 1240, the baseline is formed, with forecast data from modules 1210, 1220 and 1230.

[00129] At module 1250, changes are received, such as updates to previously forecast data. While some changes may be confirmations of forecasts (change from expectation to actual), many changes may be actual changes as orders come in at different prices, quantities and the like from forecasted data. At module 1260, a determination is made as to whether the actual changed data was watched. If so, at module 1265, notification is sent based on the watch request, with a level of detail appropriate to the request for the watch and system capabilities. At module 1270, the changes are propagated through the system, such as through an olap cube and/or database. At module 1275, the various datapoints that are affected are checked to determine if any of them are watched. If so, at module 1280, notification is provided based on the watch request. At module 1290, impacts to the forecast prior to receipt of the changed data are shown, preferably in a manner allowing for easy user interpretation of the data. At module 1295, judgment or adjustments are received from users responsive

to the changed information, resulting in the receipt of further changes at module 1250 and so forth.

[00130] Further aspects and features of an embodiment may be understood with reference to a user interface for the embodiment and a description of how and why the interface changes. Fig. 13 illustrates display by product of information in an embodiment. As mentioned, information may be displayed by region, for example. As illustrated, display 310 provides information by product, with similar or the same information. The display 360 indicates that this is a display by product, for all families in the illustrated example. Just as the impacts may be provided in impact display 315 for regions, they may also be provided for products, such that impacts may be displayed based on the view of the data provided in display 310. Alternatively, a user may specify that impacts are provided at a global level, or other predetermined/preselected level for that user for example.

[00131] Fig. 14 illustrates display by customer of information in an embodiment. Again, the option of displaying information in a different way is illustrated. Rather than display by either part or region, display by customer allows for a determination of which customers are increasing orders, decreasing orders, or experiencing short-term changes for example. As the data is preferably stored with associations to customers, parts, regions, sales persons, and other relevant information, displaying the data in various different ways becomes possible. Moreover, as the data populates an olap cube, shifting between displays involves simply shifting from displaying one aspect of the information as it is maintained in the olap cube to displaying

a different aspect of the information as it is already maintained in the olap cube. Note that the data may be displayed responsive to selecting the tab 385 (marked customers) and that display 360 indicates how the data is viewed.

[00132] As illustrated above, the data in fig. 13 and fig. 14 is displayed before a forecast is compiled into a baseline. However, changes in data may occur at any time, particularly after the baseline is formed. Fig. 15 illustrates display of a specific impact in an embodiment. Impacts are translations into human understandable language (such as english for example) of changes in data as propagated through the system. In the instance illustrated, impacts display 315 is expanded to allow for closer examination of the impacts reported. Moreover, one may drill down into the impacts, to determine how the impacts came about or what the source of the impact is. Additionally, impacts may occur before, during or after formation of a baseline forecast. Thus, status indicator 340 indicates these impacts are being viewed after a forecast has been accepted (or after the user no longer has options for input), and further indicates that continuous updates will be received.

[00133] Other presentations of the data are also available. Fig. 16 illustrates an operations report for information in an embodiment. Operations display 405 provides a view of the data based on an operational or manufacturing viewpoint, with information about changes in demand or absolute demand quantities illustrated and a summary of changes as well. Again, drilling down may also be an option, such that demand for a part may be broken down into when the demand will occur or why the changes are occurring.

[00134] Fig. 17 illustrates a forecasted change to inventory report for information in an embodiment. Unlike the view in operations display 405 of fig. 16, the inventory report illustrated in operations display 405 focuses more on monetary value of the change impact of forecast updates, with both a summary and detailed information. Such a report may also be provided for services based on capacity to fulfill demand and actual demand.

[00135] Likewise, a gap report may also be provided in operations display 405. Fig. 18 illustrates a gap report for information in an embodiment. The gap report may provide a view of variances between a projection and current actual numbers, and thus indicate what orders need to happen to ensure that the company delivers on its expected forecast.

[00136] Drilling down is typically available, unless a user is restricted from such an action. Fig. 19 illustrates a drilled down display of information for a sales person in one embodiment. Based on an earlier projection, data for a single sales person may be displayed by drilling down on a display of data for multiple sales people (such as was found in fig. 3 for example). Drill-down display 410 provides the details of the data making up the entry for one sales person. Moreover, if the user has authority to do so, apply judgment button 415 allows for entry of judgment about whether the forecast is accurate, or how it needs to change. This will be described further below.

[00137] Updates may occur while someone is viewing data such as drilled-down data, or when a user is offline. Fig. 20 illustrates an updated forecast in an embodiment. Note that changes may have occurred in the forecast for various reasons.

The display 310 provides summary data, with some of the data flagged by balls indicating changes in the forecast. Such changes may be the result of applications of judgment (post-forecast) or of actual changes to forecasted data based on real-time updates. Impacts display 315 may provide further information in some instances, such as when an actual change provides an impact.

[00138] In pursuing information about updates, one may view drilled-down data and corresponding changes or information. Fig. 21 illustrates a drilled down display of a region in an embodiment. Drill-down display 425 presents data specific to the region selected in this illustration, such as sub-region-specific, part-specific and customer-specific information. Thus, drill-down display 425 may present a different view from that provided in display 310 for example. Moreover, add watch button 420 allows a user to watch a given data point for changes, over a certain threshold, within a certain time, or otherwise, for example.

[00139] Both judgments and watches may be understood with reference to further illustrations. Fig. 22 illustrates application of judgment in an embodiment. Judgment interface 430 includes value change option 440 and percentage change option 460. Value change option 440 includes type 445 (units or revenue for example) and value 450 (amount of change for example). Percentage change option 460 includes direction 465 (increase/decrease), type 470 (revenue or units for example) and amount 475 (the amount of change for example). As illustrated, a change by value is executed, with a set amount provided – the end result of the change is specified. With a change by percentage, the amount of change (delta) is specified. Additionally, because a change

(judgment) is being entered, status 340 indicates that an approval must be provided for the change to take effect.

[00140] Fig. 23 further illustrates application of judgment in an embodiment. In this instance, since a judgment has already been applied, judgment history 480 is populated with an entry which is displayed. This may allow a user to avoid a redundant judgment entry, or may provide context for future applications of judgment. Again, a judgment is being entered in value change option 440, thus further changing the data in question.

[00141] Fig. 24 illustrates a forecast after application of judgment in an embodiment. With the change entered, a view of the original data provides an indication that judgment has been applied. In particular, judgment indicator balls 510 are now present, indicating the presence of judgment history associated with the displayed data or with data aggregated into the displayed data. As one may expect, drilling down to the changed data will allow one to see a judgment history such as history 480 of fig. 23. Fig. 25 illustrates change history for information in an embodiment. By drilling down, or by selecting the judgment indicator ball 510, the history 520 is displayed, including information about who entered the judgment, type of change, who it affected, and the amount of the change. In the illustrated example, the sales person responded to the judgment after the forecast was made by providing a further correction.

[00142] While judgments change the forecast, watches indicate changes, either actual or forecast. Fig. 26 illustrates addition of a watch in an embodiment. Watch window 530 pops up when application of a watch is selected, and includes change type

535 in which the amount of a change to be watched for is entered. Thus, the user can control whether a very small change, a catastrophic change, or some intermediate change triggers the watch. Typically, a watch sends an email or similar message to the user setting the watch once the predetermined threshold is reached or exceeded (missed or undershot).

[00143] When data is changed from an original value to a new value, such as in response to a request to apply judgment, this may have a number of effects. This may be a change in percentage or absolute terms for example, and may result from specific expectations or information, or from general expectations (hunch, intuition, etc.) For example. The change to the data results in a propagated change to other data. Alternatively, a change could be made to other data with changes back-propagated to some or all of the data contributing to the changed data. Moreover, if a piece of data is set for watching, with notification to the user should the data change, and a change propagates through, this may trigger the watch alert. Watching may be set for any change, changes above a specified threshold, or changes before a certain date for example. Also, watches may not be required to detect changes that would be visible, such that if the display generally is in billions, a change in the thousands may be sufficient to trigger a notification.

[00144] While illustrations of changes have been provided on a high-level basis (the whole world for example), changes of unit quantities may also occur. Information for a single product may be displayed over a geographical area, rather than for a geographical area for all products for example. Information may also be displayed

for a set of customers for a single product or product line, for example, or on other bases. Moreover, the information may be displayed in units of product rather than currency. Preferably, information on both a currency number and a units (quantity) number is stored. Thus, display of information in the user interface may be shifted between the two types of displays. The units may be a physical quantity (number of parts or devices for example), an estimated physical quantity (number of meals served for example), or some form of service metric (number of hours billed for example). Moreover, the various datapoints may be absolute numbers or scaled (such as quantity in thousands for example). Also, watches and judgments may be applied on a quantity basis rather than a currency basis.

[00145] Because specific products or services may be tracked, updates may be based on changes in a single order or a long-term relationship for example. A customer may decide to exit a business or discontinue a product, thus ending a need for a purchased component for example. Similarly, a vendor may decide to discontinue a product, thus requiring a customer to ramp up purchase to ensure an adequate stockpile after the component is discontinued. Such resulting updates may be a confirmation or cancellation or other change to a forecasted order, for example. Moreover, such updates may propagate further up. Additionally, any judgments may be dynamic (reduce a number by 10% always for example), or may be conditional (reduce a forecast number to x until it decreases to x). Thus, propagation may stop prior to reaching every related data point.

[00146] Further considerations for financial embodiments

[00147] The following description of an exemplary system related specifically to financial data provides details of an embodiment, along with implementation details which may be incorporated in various embodiments. The features and details may be used in part in other embodiments within the spirit and scope of the present invention, and may be combined with other features and details described previously. In particular, most of the details provided are appropriate for many types of data, and are not restricted to financial data.

[00148] The system, in one embodiment, lets companies streamline the process of creating a bottom-up forecast of sales and financial data. In one embodiment, this includes collecting the data from those on the front lines. This may include receiving data from sales representatives, distributors, representative firms, customers, retailers, or other sources of forecast data. Following this, the data may be aggregated within the system, and the system may then allow sales and marketing management to apply judgment. The hierarchical judgment applications may be tracked, such as by maintaining data about judgments applied to specific data and to corresponding changes to other data. Moreover, prior to or after applying judgment, the consequences of the judgment applied may be understood, as the changes flow through the data (as implemented by the olap cube), and are displayed. Moreover, analytical tools may also be employed to understand the data. Examples of such tools include regression analysis, statistical analysis, data mining, and correlation analysis among other tools.

[00149] After an initial review of the data provided, a sales vice president or similar person in authority may create an official forecast baseline, preferably after the person has understood, judged and approved the data that has been rolled up to him. At this point, the person may release the data for others in the company to consume; and define the baseline against which updates will be tracked. The baseline may be yearly, quarterly, monthly, bi-weekly, weekly, even daily if desired, and may be implemented on some other time-frame. Moreover, multiple users or persons in authority may play a role in building the forecast or modifying the forecast, such as by allowing for marketing input for example. Thus, some portions of the forecast may come from marketing; marketing may apply judgment at some point in the process; or marketing may provide longer-term forecasting (in contrast to shorter-term sales forecasting) for example.

[00150] The released forecast then provides departments within the company, such as production, engineering and operations, insight into what will need to be built, both when and where. Moreover, this allow for vetting (and thus feedback) from production, and may allow for prediction of trends for parts or supplies for example. Similarly, this allows finance departments to analyze and predict financial data such as a gross margin, either on a line-by-line level or at an enterprise level, for example. This then allows for planning of capital needs and for simulation or 'what if' type of scenarios, both within the system or in a separate financial system. Additionally, finance departments can provide feedback to the baseline as well, such as by indicating which accounts are doubtful and should be discouraged until payment is more reasonable, or by indicating what expected financial trends may do to various industries.

[00151] From this, the ceo may then see all perspectives of the forecast, along with the broad overview of the forecast. This allows the ceo to obtain 'one number' for the entire company – allowing for intelligent discussions with media and outside interests when the ceo interacts with the public. Alternatively, the ceo can track a set of numbers and associated relationships, allowing the ceo to also understand the perspectives within the company which relate to these numbers. Moreover, this provides a clear and detailed view of expected developments of the company. As the ceo may also simulate changes, apply judgments or watch numbers (along with other departments and people), the ceo and staff may then analyze potential changes. As such, this facilitates key decisions a ceo may need to make. Such decisions, whether made by the ceo or some other member of the company, may include determining what parts to retire and when; how much of internal resources to allocate; where to invest based on what appears to be driving the business; how to streamline internal operations; and how to maximize capital efficiency for example.

[00152] With the baseline in place, not only analysis and feedback, but also real-time updates are available. The system may show the impact of changes if they pass a certain threshold, and show the impacts by department, group or otherwise. This then allows the company to react to external (market for example) forces, allows all groups to consider and agree on options to handle changes, and allows for a group or consensus decision on whether or not to choose specific options.

[00153] In one embodiment, the process may be described as follows:

[00154] Data collection is automated, to the extent possible. This may include sending out automated reminders to those generating the data (or inputting the data they observe for example). The users or forecasters (who may be sales people with many other demands on their time) enter data using a simple user interface which is robust enough to trap or catch common errors such as entering too many or too few zeros, unintentionally large or small changes, or incomplete data or omissions for example. Moreover, the user may be provided data to help create the right forecast. For example, a backlog may be used as a starting point, or the last forecast may be used as such a starting point. Additionally, feedback may be provided on forecast accuracy, such as by attempting to curb over-optimistic forecasts or sandbagging.

[00155] With data entered, or with a deadline approaching or past, notifications may be sent to those who are tasked with reviewing and judging the data. Notices may relate to delinquent forecasts, forecasts in and ready for approval, or problems with forecast creation for example. The system provides tools to view the data, view aggregated data (which may be viewed or drilled down to various levels), apply judgment (globally or locally for example), send forecasts back for rework/correction/reconsideration for example, and approve the forecast (and/or send to the next hierarchical level).

[00156] Preferably, the interface is simple, with controls that change global perspective (such as switching to a graphic or tabular view for example). This may further include providing comparisons to various time periods (such as a previous forecast, successive quarters, year on year comparisons, and current information versus

baseline comparisons for example). Similarly, comparisons to actual sales or actual data may be provided.

[00157] Moreover, data may be shown in various forms, or with various aspects of the overall collection of data displayed, and data may be drilled down from higher to lower levels of data aggregation, ultimately to atomic data levels. Thus, data may be provided as revenue, units, gross margin currency, gross margin percentage, simple margin currency, simple margin percentage, average selling price, or in any other format either collected or derivable from collected data. Similarly, data may be sliced by various means, such as by region, customer, part (or service), or by some custom aggregation of atomic data or previously aggregated data for example. A custom aggregation may allow for display of data by program, segment or other varying groupings of data for example.

[00158] Moreover, data may be displayed based on various breakdowns. For example, users can drill into any value at any point by clicking and may then see what makes it up. Thus, users may break down data on parts, customers, or regions for example. Similarly, users can then click on any sub-entity (such as a division of a customer or a part of a region for example) to get further breakdowns.

[00159] In some embodiments, baseline creation and continuous updates form separate tasks and may have separate interfaces. Thus, a different business process may be used for each. Creating a baseline is often oriented towards a bottom-up commit – sales people providing commitments to the company for example. Creating updates is

often more of a change notification prompting or requiring decisions and/or action – the biggest customer starts canceling orders and high-level executives need to act right away.

[00160] With the data stored in the system, this enables continuous updates – a change may be propagated essentially immediately. The system may immediately notify all groups to a change in terms that make the impact (of the change) clear. These groups may include sales, marketing, production, engineering, operations, finance, and executive groups for example. Moreover, the updates may be provided using a simple client or interface, and may be tied to production to ensure that sales people will get goods for their customers in the right quantities at the right times.

[00161] While continuous and essentially real-time updates provide a fairly accurate picture as changes occur, judgment may be used to predict changes, and judgment histories may be employed to determine when predicted changes occur, or whether predictions are already integrated into a forecast. Judgments may be applied hierarchically to forecast data. Thus, a judgment may be applied to a high-level number, and then changes may be cascaded or propagated to lower levels, allowing users to see the effects of the judgment.

[00162] Tracking and inspecting the history of judgments applied (by attaching judgments and judgment history to data for example) allows users to determine whether judgments should be reversed for example. Moreover, variances of actual results from judgments may be examined and analyzed. Additionally, as judgment may be applied prior to or after a baseline is formed, differences of pre-baseline versus post baseline judgment may be tracked.

[00163] While judgments allow for predictions, watches allow for action. Every user has the opportunity to design custom notifications around anything they can see from their vantage point in the system. Thus, the user may choose to watch a data point for a particular customer, part, service, time period, or other data point, as long as it is visible. The user may set a custom threshold for notification, and a users may create and save collections of watches with corresponding messages that can be turned into reports for example.

[00164] Further analysis and related judgments and watches may be applied to custom aggregations of data. Preferably, the system lets customers aggregate data in any way they choose. For example, part data may be aggregated based on programs. Customer data may be aggregated based on market segments, region data may be aggregated based on climates, selling entity data (salespeople, representative, etc.) May be aggregated based on performance. Each of these aggregations are exemplary, and illustrative of the options for aggregation provided in part due to use of the olap cube. These aggregations may be employed for a variety of purposes, including simple reporting, analysis, or export to other systems for example.

[00165] Impacts of changes are often an important goal of analysis, watches, judgments, aggregations or any other exercises carried out on financial data. The system automatically determines what the biggest impacts are and displays them by absolute rank and relative change percentage – basically by tracking changes as they occur and maintaining a set of lists of such information. This allows users to drill down into sources of the biggest impacts. Moreover, impacts are provided according to

whatever the drill down perspective is, (such as through context sensitive impacts) in some embodiments.

[00166] Multiple organization considerations

[00167] The system described above can be used profitably by a single organization. However, having multiple organizations using the same system requires that data be segregated to satisfy confidentiality concerns, for example. As indicated previously, data repositories may be separate, either logically or physically (or both). Moreover, processing facilities may also be separated, logically and/or physically. Thus, a first user may use a first client to access a system, and thereby access a first data repository using an olap cube. A second user may use a second client to access the system. In so doing, the second user may access a second data repository, and would preferably use a different olap cube or at least a different instantiation of an olap cube or cubes.

[00168] Thus, each user, or at least each organization, has access to olap cubes which in turn work only with data from a dedicated repository. Moreover, each dedicated repository may have associated with it customizations for the organization and/or user in question. Thus, an instance of an olap cube may effectively be customized for the user or organization when access to the dedicated repository is granted.

[00169] By providing dedicated instances of olap cubes and dedicated repositories, a flexible structure which may be hosted across multiple servers is formed. This supports providing a web-hosted application, such as through use of asp technology.

Multiple users or organizations may be supported through dedicated repositories, dedicated instances of olap cubes, and shared supporting software and physical resources. Similarly, this architecture does not tie down the location of physical resources, allowing for either distributed resources (such as geographically separated servers and networks for example) or concentrated resources (such as server farms for example). Considerations such as geographic diversity/redundancy or ease of maintenance may come into play because the technology allows for such flexibility.

[00170] Further examples of various embodiments

[00171] Specific details of various embodiments may be understood with reference to the following description and accompanying figures. The various embodiments may share features with other embodiments, such that an embodiment may incorporate features from several different embodiments illustrated herein, even though those features are not specifically illustrated together in a single embodiment. Similarly, various embodiments need not include all of the features or aspects illustrated in an embodiment, provided the embodiment still meets the claims.

[00172] Feedback

[00173] One example of a feature present in some embodiments is a feedback capability. Fig. 27 illustrates an embodiment of a process of providing feedback in a real-time forecasting system. Process 2700 relates to providing such feedback, and includes collecting data (receiving data), receiving feedback on that data,

providing the altered data or data with feedback to the original user providing the original data, and receiving and checking for approval further changes to the data after feedback is received.

[00174] Process 2700 begins at module 2710 with receipt of data, such as forecast data entered by a field operative for an enterprise. At module 2720, the entered data is integrated or submitted into a database, and may also be populated into an olap cube, for example. The data is then provided for review at module 2730, presumably with data collected from other field operatives. Feedback information is provided for the data at module 2740, including changes to data and other forms of feedback. This feedback may take on various forms, and may be entered by various people or entities. At module 2745, the feedback data is integrated or entered into the corresponding database, and may also be populated in an olap cube, for example. Modules 2730, 2740 and 2745 may form a subprocess which is repeated for various people, organizations or groups, for example.

[00175] At module 2750, a request is received for data in its current state from a field operative. Thus, the field operative is seeking data after feedback information has been provided. The altered data is provided in response to the request at module 2760, such as through a database request for information related to a given user, for example. Thus, the user may view information or data which is visible based on an associated userid (e.g. Data the user has privileges for or is authorized for) or may view data based on what information the user originally entered, for example. At module 2770, the user may provide objections to changes to data, such as contesting changes to

some or all of the data, or contesting the reasoning behind such changes if available. These objections are received by the system, and may be submitted for review by a predetermined user at module 2780 – thereby allowing for intelligent adoption or rejection of objections or changes. At module 2790, the data in question is finalized, signaling that the data will not change further, at least during the present data entry/modification cycle.

[00176] In some embodiments, real-time forecasting may be used in commercial enterprises to forecast sales, production, and related revenues, for example. Fig. 28 illustrates relationships between various groups and data associated with those groups in an embodiment. System 2800 represents the internal system for forecasting of a commercial enterprise, with relationships between various departments illustrated in an exemplary manner.

[00177] Sales group 2810 may present an initial sales forecast 2815, representing an expectation of what sales will be for an upcoming period of time, such as an upcoming month. This sales forecast 2815 may provide a baseline for comparison against actual performance and against forecasts from other parts of the organization. Based on sales forecast 2815, marketing group 2820 may then provide a marketing forecast 2825, providing input on how marketing thinks actual results will differ from the sales forecast 2815, and potentially providing reasoning for differences as well. This input may be provided on a macro or micro level, either as a broad adjustment to an entire category (geographical region, customer, industry, product, for example) or to

specific forecasts from particular people, for example. Thus, forecast 2825 may form a marketing baseline in some instances.

[00178] Building on forecast 2825, finance group 2830 may then provide forecast 2835, which may be expected to reflect finance group 2830 judgment about forecast 2825. The judgment reflected in forecast 2835 may reflect expectations about cash flow, customer payments, inventory, or other areas of concern tracked by finance group 2830. Similarly, operations group 2840 may provide forecast 2845 based on forecast 2835 and expectations about capacity or other operational factors. Finally, executive group 2850 may use forecast 2845 as a baseline to generate forecast 2855, after applying expectations of executive group 2850 to the data of forecast 2845. Note that forecast 2855 may be dependent on components of various forecasts or baselines, such that the finance baseline 2835 may be deemed more reliable, or the marketing baseline 2825 may be viewed as trustworthy by executive group 2850. Similarly, this strategy may be used for parts of forecasts, thus allowing a baseline 2845 to be adopted for some areas, but allowing other areas of an ultimate forecast to start from one of baselines 2815, 2825 or 2835, for example.

[00179] Preparation of baselines and achievement of related objectives involves forecasting expectations, providing feedback on those forecasts, and reviewing objections to feedback. Fig. 29 illustrates another embodiment of a process of providing feedback in a real-time forecasting system. Process 2900 includes gathering data, presenting the data, gathering feedback, presenting the feedback, gathering objections to the feedback, and reviewing objections.

[00180] Process 2900 begins with gathering of forecast data at module 2910. Forecast data is then presented to reviewers at module 2920. Feedback (changes, commentary, for example) on the forecast data is gathered from reviewers at module 2930. The feedback is then presented to providers of the forecast data at module 2940. Objections to the feedback are then gathered at module 2950, and those objections are reviewed as appropriate at module 2960.

[00181] Where data comes from (and where it goes) in a forecasting system may illustrate how feedback is provided. Fig. 30 illustrates a real-time forecasting system and its relationship with various groups in an embodiment. System 3000 includes a forecasting system, and interfaces (interaction) with sales, marketing, finance, operations and executive groups. The information flowing back and forth between the forecasting system and the various groups provides for forecasting, feedback and changes to forecasts, and training on future forecasts.

[00182] Forecast system 3060 includes a database 3070 which embodies forecast data and an olap engine 3080 which interacts with database 3070 to provide access to data in particularly flexible ways. Sales group 3010 provides forecast data to system 3060 and receives back forecast feedback information, such as indications of whether forecast data seems unreasonable, has been changed due to market conditions, or has been changed due to prior forecast performance, for example. Sales group 3010 may also provide commentary or objections to feedback, either due to organizational commitments or due to differing judgments, for example.

[00183] Marketing group 3030 similarly reviews sales forecasts and provides overall market data, either through specific data input or through changes and comments on sales forecast data. Finance group 3040 receives/reviews an overall forecast, and provides information on financial constraints, such as through separate data inputs or through suggested changes to forecast data, for example. Similarly, operations group 3050 reviews the current forecast data and provides data on production constraints, such as by adjusting forecast data for example.

[00184] All of this data interaction provides a forecast which executive group 3020 may review and alter. Alteration may occur through application of executive judgment to the forecast, based on knowledge of market conditions, separate conversations with customers, or other sources of knowledge. Such alteration may include commentary along with changes, and the act of alteration may be recorded to allow a user to understand how changes occurred. Thus, sales group 3010 may provide forecast sales data, marketing group 3030, operations group 3050 and finance group 3040 may comment on and change the data, and executive group 3020 may also review, alter and comment on the forecast data, to arrive at an overall forecast. Moreover, various groups may then review the overall forecast to determine how individual forecasts were changed or commented on by other groups, allowing for improved future forecasting and for understanding of expected changes in forecasts.

[00185] A specific process for sales forecasting and feedback may further illustrate the forecasting system. Fig. 31 illustrates an embodiment of a process of providing sales feedback in a real-time forecasting system. Process 3100 includes

receiving sales data, storing the data, presenting data for review, receiving and integrating feedback, receiving a request from a salesperson, presenting altered data, receiving objections and updates, presenting objections and updates for approval, receiving such approval (or denial) and integrating data into the database.

[00186] Sales data in the form of forecasts for upcoming time periods is provided at module 3110. This data is stored in a database (as new data or changes) at module 3115. The data is then presented to other stakeholders for review at module 3120. Such other stakeholders may include marketing, production/operations, finance and executive groups, for example. Feedback on the forecast data is provided at module 3130. This feedback, in the form of changes to data, comments on data, and similar changes, is integrated into the database at module 3135. A determination is then made at module 3140 as to whether all feedback has been received. This may relate to receiving feedback from all departments, or may relate to time for feedback ending, for example. If feedback is still to be received, the process returns to module 3120.

[00187] If all feedback is received, the process then moves to providing feedback to users. A request from a salesperson for feedback is received at module 3150. This request may be an affirmative request from the salesperson, or may simply be an indication that the salesperson has logged into a system at a time when feedback is automatically provided. Altered data or data with commentary is presented at module 3160, such as by allowing a salesperson to see what changes were made or see what impacts changes had. If the salesperson has objections to changes or updates to data, such objections and updates are received at module 3170. As these objections and

updates (changes) are provided after the forecasting period, these changes may be reviewed at module 3180 prior to integration into the forecast. Those changes that are allowed are approved at module 3190, and other changes are denied. The approved changes are then integrated into the database at module 3195.

[00188] Understanding potential changes and related feedback of an atomic data point may be useful. Fig. 32 illustrates changes to an atomic data point in one embodiment. Chronology 3200 illustrates changes to such a data point. Data at point 3210 includes a forecast of 400 units for xyz company in the northeast. Point 3220 indicates feedback of a softer market (less demand as seen by a marketing department for example) indicates only 350 units will be ordered. Point 3230 has no changes and no opinion expressed (such as a finance department seeing no need for change).

[00189] Point 3240 illustrates a further downgrade in an order (or units to fill the order) due to a production constraint (as may come from an operation or production group for example). Point 3250 includes an upgrade in the forecast, back to 350 units, indicating capacity should be found (such as when an executive decision about satisfying a customer occurs for example). Point 3260 illustrates feedback which may be seen by the sales person entering the data, showing the various changes and related commentary.

[00190] Note that each of the points in the chronology may also support a baseline, and some or all data points may be used as an objective to be achieved. The baseline may be generated due to timing of data entry (meeting a deadline) or due to a desire to capture a state of the system (taking a snapshot for later reference). Similarly,

an objective may be related to an organizational interest in meeting commitments, for example.

[00191] Baselines

[00192] Generation of a baseline from which objectives may be measured can occur in a variety of ways. Fig. 33 illustrates an embodiment of a process of generating a baseline. Process 3300 includes receiving data, integrating the data into a database, generating a snapshot, and classifying the snapshot as a specific baseline.

[00193] Process 3300 includes receiving data at module 3310. The data received may be original data, or a change to existing data. If the data is original, it will be included in the database at module 3320 as a new data entry. If the data is a change, it will be included in the database along with context information about what the change is. Moreover, the data may be a simple numerical value or change, or it may be a comment or annotation on an existing value, for example. All data received includes a timestamp within the database, indicating when it was entered. Receiving and integrating data continues until a cutoff point is reached.

[00194] A determination of a cutoff is made at module 3330, and may occur in a variety of ways. Some cutoffs may be automatic, such as generation of a baseline on a preset day of the current week or month. Other cutoffs may be manual, such as a request for a baseline from a user or department. When the cutoff occurs, a snapshot is generated at module 3340.

[00195] The snapshot may be implemented in a variety of ways, too. For example, a snapshot may be a recording of the time at which the baseline was requested, or a preset time for a baseline previously requested. This snapshot time marks which changes should be included to calculate a baseline, and which changes should be excluded as they occurred after the time of the baseline. Alternatively, a snapshot of the state of the system may be recorded, much like a backup of data in a static fashion. Such a backup may be resource and time-intensive to complete, but it may also facilitate calculation of differences between values in the baseline or values of the baseline and other values later.

[00196] The snapshot is classified as a baseline at module 3350. This may simply include assigning an identifier to the value of the snapshot (or the data stored as the snapshot). Additionally, it may include designation of permissions related to users authorized to view a baseline. For example, a baseline may be a private baseline, created for personal use of a user or restricted use of a group. Alternatively, a baseline may be restricted to those having certain general access rights, although such access rights may change over time, either in an evolutionary manner with an organization or in a cyclical manner, for example. Additionally, the baseline may be restricted in what it includes, such as by restricting it based on various categories within the data such as customer, data source (marketing and operations but not sales, for example), or by restricting it based on how far back data may be retrieved to provide the baseline.

[00197] Baselines may be used in a variety of ways. Fig. 34 illustrates an embodiment of a process of operating a real-time forecasting system in conjunction with

a baseline. Process 3400 includes receiving and integrating data, comparing data to established baselines, checking for cutoff points for generation of baselines, and generation of baselines.

[00198] Data is received at module 3410, and is then integrated into a database at module 3420. This integration may also include integration into an olap server and data structure, as will be described in greater detail below. If a comparison is desired (requested by a user) then a determination is made to perform a comparison at module 3430. Such a comparison occurs at module 3440, using a baseline previously established as one of baselines 3450. Thus, a user may determine how much a current forecast as represented by data in the database has varied from a baseline forecast established at a previous time. Alternatively, the user may determine how much two different baselines diverged, allowing for post-mortem analysis of changes in data made by various users, and for analysis of differences between forecasts and results.

[00199] At module 3460, a determination is made as to whether a cutoff point has been reached. If not, the process returns to receiving data at module 3410. If so, such as due to a predetermined cutoff or a request for a baseline, a snapshot is generated at module 3470. This snapshot may be generated as it would be for module 3340 of fig. 33. The information of the baseline is then classified as a baseline at module 3480 and added to baselines 3450 for future use in comparisons. Note that such classification may be similar to that of module 3350 of fig. 33.

[00200] Also, note that baselines may, and typically will, be derived from earlier baselines. Typically, baselines will derive in a sequential fashion, one from the

previous one. However, private baselines may be derived in some instances, and such baselines may be derived from a selected subset of past data as incorporated in various baselines. The windows discussed later may be used for this purpose, with only some of the earlier windows specified for a particular private window or baseline.

[00201] Baselines may provide a variety of opportunities for analysis. Fig. 35 illustrates relationships between baselines and potential operations in an embodiment. Data collection 3500 includes multiple baselines and illustrates results of analysis that may be achieved. Baseline 3510 is an original baseline in a forecast, such as a baseline established after sales data is initially received. Revised baseline 3520 is a later revised baseline, which may include input from a variety of sources, and potentially has different data from that in baseline 3510. Actual results 3530 are the actual results of operations corresponding to the forecasts of baselines 3510 and 3520.

[00202] Comparison of this data results in various analytical results. Variance 3540 represents differences between two sets of data, and illustrates where the forecast did not meet reality. Accuracy 3550 represents how far away a forecast was from reality or a later forecast – accuracy 3550 and variance 3540 may represent the same data, and either present the data differently, or present different aspects of the data, such as percentage versus absolute magnitude for example.

[00203] Distortions 3560 represents identified factors distorting the forecast, such as market changes, customer defections, unpredictable events, and other issues which may account for differences between a prediction and a result. Tendency 3570 represents identified trends in predictions, as may be represented by tracking

forecasts over time or tracking results of forecasts by users over time. Tendency 3570 may include identification of forecasters who are unpredictable (unreliably inaccurate), regularly optimistic, regularly pessimistic, or generally accurate, for example. Similarly, tendency 3570 may illustrate correlations between different forecasts, or between forecasts and results, for example. Correction 3580 may include an identification of factors which may be used to correct future forecasts. This may include identifying users who should be trained to forecast better. It may also include identifying offsets or multipliers which may be used to correct forecasts of those who are reliably incorrect, for example.

[00204] Annotations

[00205] While baselines allow for much comparison of results, annotation of results reinforces or potentially corrects predictions during the predictive process. Moreover, annotation of data may raise issues (either actual or anticipated) at an early stage to allow for discussion and resolution of such issues. Annotations may be gathered after initial data is gathered, to indicate whether the initial data looks incorrect to users of the system, and allow for more meaningful interaction between users that simple changes to numbers allows.

[00206] Thus, the data and annotation gathering process may include gathering data, annotating data, displaying annotated data, and revising annotated data. Fig. 36 illustrates an embodiment of a process of gathering data in a real-time forecasting

system. Process 3600 includes receiving forecast data, integrating data into the database, determining data gathering should stop, and halting gathering of data.

[00207] Forecast data is received at module 3610. This may include initial forecast data from various users such as field operators or sales people, and may relate to what is expected in various timeframes. Typically, the forecast data is a personal forecast which includes business information which is expected to be useful to the organization. The forecast data is integrated into a database 3650 at module 3620. This integration may be entry of initial data with associated parameters, or may be entry of a change to a prior forecast in the case of longer term forecasts. At some point, a cutoff point may be reached, such as an end of a forecasting period. This is determined at module 3630, and when the cutoff is reached, the data entry process halts at module 3640, such as through locking of database 3650 to further changes or changes from designated users. Note that database 3650 may include a classical database, and an associated olap structure as well.

[00208] With initial forecast data entered, the forecast is then reviewed by other users. Fig. 37 illustrates an embodiment of a process of annotating data in a real-time forecasting system. Process 3700 includes presenting data to users, receiving adjustments and comments, integrating adjustments and comments into the data base, and completing annotation entry.

[00209] At module 3710, data from database 3650 is presented to a user or users. This typically occurs on a wide scale, allowing various users to review and change or comment on data. Adjustments to data are received at module 3720, and comments are received at module 3730. Note that such adjustments (actual changes) and comments

(notes about the data) may be received essentially simultaneously, and for a variety of different data points. At module 3740, the adjustments and comments are integrated into database 3650. If the annotation process is complete, this is determined at module 3750. If it is not, then the process returns to module 3710 for presentation of data, whereas if the process is complete, module 3760 halts entry of annotations.

[00210] In some embodiments, the initial forecast data is compiled from field operatives such as sales people. Annotations then come from marketing groups, production/operations groups, finance groups and executive groups. Thus, the process of modules 3710 through 3750 may be repeated several times for different groups or users, even for the same piece of atomic data. Also, later users may be able to see some or all annotations from earlier users, allowing for further commentary or avoiding duplicative (and overly cumulative) changes.

[00211] After entry of annotations, the data may be reviewed by users who originally entered it, along with the annotations. Fig. 38 illustrates an embodiment of a process of displaying annotated data in a real-time forecasting system. Process 3800 includes receiving a request for annotated data, retrieving the requested data, and presenting the data to the requestor.

[00212] At module 3810, a request for annotated data is received from a user. For example, a sales person may wish to see what marketing, production and finance did with previous projections and whether desired commitments will be met. Alternatively, a sales person may want to see what feedback is available about the overall market or the overall customer relationship. At module 3820, requested data is retrieved

from database 3650. Note that requested data is likely to be data requested by the user through a user attempt to access the data, rather than a specific request from the user. At module 3830, the requested data, now retrieved from database 3650 along with associated annotations, is provided to the user, such as through a graphical user interface. Note that if the user has limited rights to access the data, some annotations may not be available to the user in question.

[00213] The nature of human interaction is that disagreements arise. Fig. 39 illustrates an embodiment of a process of revising annotated data in a real-time forecasting system. A sales person may disagree with changes to forecasts, either because specific information contradicts a more general change or because the sales person simply believes the data to be inaccurate. Both situations may be addressed.

[00214] Process 3900 includes receiving comments from data reviewed, integrating comments into the database, and finalizing comments in the database. A sales person or similar forecaster may provide comments about data presented with annotations at module 3910. Such comments may contest annotations and changes, request further changes, or indicate changes are agreeable, for example. At module 3920, these comments are integrated into the database 3650. This integration process may simply be recording the comments, or may involve an effort to alert a supervisory user to the comments. At module 3930, comments in the database are finalized. This may include review by a supervisory user, review by those affected by comments, or some other form of check on the reason for changes.

[00215] Note that annotations may involve more active functions than simply providing an opinion or change. For example, in some embodiments, annotations not only raise an issue through a comment, but also provide an indication of a lack of confidence in part or all of a data point without specifically changing that data point. For example, if a sales person indicates 500 parts will be ordered by a customer six months later, that can be annotated to indicate it is contingent on the customer designing in the specified part in a new product.

[00216] Thus, the annotation may not only state the concern, but include a tag indicating how much the number should be reduced to handle the contingency. If the part would be ordered in a quantity of 300 units without the design-in, the annotation may indicate 200 units are contingent on a design-in decision. If the decision is slated for a month after the annotation is made, the annotation may be implemented to notify the user making the annotation, and the user whose forecast was annotated, to check on the decision. Such notifications may be email messages sent at specified or predetermined times, and may also show up along with impacts in a dashboard display, for example.

[00217] The overall process of data collection, annotation, and further revision, may be even more involved in some embodiments. Fig. 40 illustrates an embodiment of a process of operating a real-time forecasting system. Process 4000 illustrates such collection of data, collection of changes and annotations, collection of objections, and use of the data in both forecasting and review of execution. The actual implementation of this process and similar processes may vary in terms of order of operations, use of all modules, or addition of some modules, for example.

[00218] A sales forecast is received from a user or set of users at module 4010. This sales forecast may be initial data related to expected sales, for example. At module 4015, the sales forecast is revised. This may include revisions by those originally entering the data, or revisions by others in the sales group, such as superiors. At module 4020, a baseline is generated using the sales forecasts from the sales group/department.

[00219] At module 4025, marketing group updates to the forecast are received. These may be annotations, either comments or changes, entered by various marketing personnel. Such changes may relate to market conditions, overall relationships with customers, and later arriving data, for example. At module 4030, a finance group update is similarly received. The finance group may provide information related to financing constraints or financial market conditions which impact sales forecasts. In a similar manner, an operations or productions set of inputs is received at module 4035. This may include production constraints, inventory considerations, other supply chain issues, and allocations of production capacity, for example. All of these inputs may then be used to generate a new baseline at module 4040. The inputs of modules 4025, 4030 and 4035 may include judgments, watches, changes and comments, for example.

[00220] At module 4050, the forecast undergoes executive review. This may involve a review of high-level information (company-wide or organization-wide performance), and may further include drilling down to more specific information in some parts of the forecast. Moreover, executive updates to the forecast may be provided at module 4055. Such updates may reflect expectations for the economy, customer relationships and overall business, or may reflect executive expectations about actual

performance (as opposed to forecasted performance) by various parts of the company. Entry of information at module 4055 may include judgments and watches, along with other annotations. At module 4060, another baseline is generated, incorporating executive changes to the forecast.

[00221] Feedback is then provided to stakeholders, such as those originally entering information and those entering later changes, at module 4065. This feedback may include annotations and associated updates to forecasts. Note, during the process of entering data, watches may have alerted people to some changes already. Module 4065 may implement a more formal or thorough notification. Objections to changes may be received from stakeholders at module 4070. Such objections may relate to differences of opinions, or to a need to meet organizational and personal commitments to clients or customers, for example. Similarly, revisions to the forecast may be received at module 4075. Objections and revisions may require authorization or approval for actual entry, or may be entered without intervention in various embodiments. Yet another baseline may then be generated at module 4080.

[00222] Actual performance data may then be collected at module 4085. Thus, actual sales figures or performance information may be collected. At module 4090, actual performance information may then be compared to the various baselines. Moreover, reports of differences between execution (actual data) and forecasts may be generated for use by various users.

[00223] Note that comparison is illustrated at the end of the process. Comparison of various baselines during the process may also be useful. Moreover,

determination of why numbers in various baselines do not match may allow for avoidance of hidden problems, for example. Thus, the process may be manipulated by users in various embodiments to achieve useful business results.

[00224] Reviewing data collection in a timeline model may also be illustrative. Fig. 41 illustrates a timeline of operations in conjunction with a real-time forecasting system in one embodiment. Timeline 4100 illustrates data gathering and baseline generation in the process in one embodiment. Sales forecast data gathering occurs, and is exemplified by data collection events 4110 and 4115. These may be deadlines for forecast reporting, for example. Baseline generation event 4120 may mark the end of sale forecast collection.

[00225] Internal review may include review by marketing, finance and operations groups, for example. Data collection events 4125, 4165 and 4175 illustrate various collection events and data entry by users or groups of users. Some data collection may involve group meetings to decide on what data is entered, other data collection may be individual. A baseline is generated at event 4130, which may be an internal baseline, for example.

[00226] Executive review may involve further data collection at events 4135 and 4160. Such data may be collected based on developments in the business, meetings with clients, or other events. An executive baseline is the generated at event 4140. Execution then occurs, although some execution may have been occurring on an ongoing basis (resulting in business developments for example). Actual data is gathered

at events 4145, 4170 and 4150, and an actual baseline is generated at event 4150 to allow for comparisons to the prior baselines.

[00227] Further technical features

[00228] Feedback, baselines and annotations may all be supported by various technologies. Fig. 42 illustrates an embodiment of a data structure. Data structure 4200 represents data in a data base, with various tables or dimensions of the data illustrated. Other data may also be involved, depending on implementations and database schema in various embodiments.

[00229] Fact data 4210 is the actual data collected in a system. This may be facts such as number of units to be shipped to a customer, for example. User data 4230 is related data indicating where the fact came from. Customer data 4230 is related data indicating which customer the fact datum of 4210 relates to. Product data 4240 is related data indicating which product (or service) the fact datum of 4210 relates to. Thus, a sales person may enter the number 400, representing a number of units or price of units or similar information. From the user interface (where the data was entered), the customer and product may be determined. Similarly, the identity of the user may come from a login procedure. Additional data such as the time of the entry, the timing of the data (when the goods will be shipped or service will be rendered), the type of entry (units, currency, volume, etc.) May be collected and stored as well.

[00230] Accompanying the database may be an olap structure as discussed previously. Fig. 43 illustrates an embodiment of an olap structure. The olap system 4300

is illustrated with two partitions, though the number of partitions may vary based on various implementation details. In this embodiment, a first partition 4310 includes general data from a database. This partition is updated periodically, such as once a day for example. A second partition 4320 includes change data. This partition is updated much more frequently than partition 4310, and may receive data from a database on a continuous basis, for example.

[00231] Updating the olap structure includes both adding data and calculating the effects of that data. Thus, partition 4320, containing change data, may be calculated frequently without using too many resources due to the small amount of data included, whereas partition 4310 with much more data may be calculated and added to less frequently. Changes 4330 are sent to partition 4320, such as from a database. This system may perform particularly well in response to a query 4340. Such a query may be sent to both partitions, receiving data from both, and having the data aggregated to appear as a single response.

[00232] Making the changes in an olap structure is part of the process of providing a fast query response. Fig. 44 illustrates an embodiment of a process of making changes in an olap structure. Process 4400 includes receiving data, entering the data in a database with a flag set, processing the change partition, determining if a periodic update should occur, processing all partitions, clearing flags, and using a unified partition.

[00233] A fact datum (new data point or change to data) is received at module 4410. At module 4420, the datum is recorded in a database, with a flag

indicating a change has occurred – the recorded datum is new. At module 4430, all data in the database which is flagged is then processed into the change partition of an olap cube. Since this is likely to be a small amount of data relative to the contents of the database, this results in relatively low consumption of resources. Thus, the olap cube has a partition with changed data in it, and that data populates the cube to allow for access to information in a speedy fashion.

[00234] At module 4440, a determination is made as to whether the entire olap cube should be updated. This may occur in conjunction with a timer 4445, which may be set to alert to a need for update on a daily or other predetermined basis, for example. If no update is due, the process returns to module 4410 for the next change. If an update is due, changed data is provided to all partitions, and all partitions are processed at module 4450. All flags related to changed data are also cleared at module 4460. This results in a unified partition at module 4470, which may be used to answer queries. As changes accumulate, queries will tend to invoke both the change partition and the previously unified partition, through the same process 4400.

[00235] With the data in the olap structure, queries may be sent and responsive data supplied. Fig. 45 illustrates an embodiment of a process of responding to a query in an olap structure. Process 4500 includes receiving a query, sending the query to both cubes, receiving data, aggregating the data, combining the results, and providing a user result.

[00236] A query is received at module 4510, requesting data from the olap cube and associated database – and which will be serviced from the olap cube.

Preferably, the query is then sent in parallel to the two cubes (partitions) or in a nearly parallel manner. At module 4520, the query is passed to the main cube. At module 4530, responsive data is received from the main cube, and at module 4540, that data is aggregated into a result.

[00237] Correspondingly, at module 4545, a determination is made as to whether any changes are in the change cube. If not, the response of module 4540 is complete at module 4580. If changes are present, the query is sent to the change cube at module 4550. At module 4560, data is received from the change cube responsive to the query. At module 4570, that data is aggregated, and at module 4580, data from the two cubes is combined into a single response. At module 4590 that response is then provided to the user. Aggregation may involve computing final data or combining parts of data which were not previously combined with the cube, or organizing data from a cube into a more useful format, for example.

[00238] Data from the cubes and database may be organized based on baselines. Fig. 46 illustrates an embodiment of a data structure in conjunction with a timeline. Data structure 4600 is illustrated in conjunction with timeline 4650 to illustrate how baseline information may be stored.

[00239] Data structure 4600 includes facts 4610, customer information 4620, user information 4630, and other associated information. Data structure 4600 also includes window information 4640. Thus, for each fact datum in facts 4610, a customer, user, and window is associated – although this data need not be absolutely populated for all data.

[00240] Timeline 4650 illustrates a timeline for a monthly cycle or baselines. A first baseline or field baseline is set at day 7, and designated as window #1. A second baseline is set at day 15, and designated as window #2. A third baseline is set at day 20, and designated as window #3. For each of these windows, data entered prior to the baseline is included and data entered after is excluded. Thus, a fact datum may have a window number associated with it, indicating when it was entered (along with a timestamp which is not shown). Alternatively, timestamps of data may be used in baseline calculation, and windows 4640 may store the set of baselines and associated timestamps for comparisons.

[00241] Windows may also be designated privately, thus indicating that additional changes were made in a what-if type of scenario and should not be retained for all forecasts, but should be used in a separate baseline for purposes of experimentation with forecasts. Thus, a window 21 may be designated, for example, with changes associated therewith, and that window 21 and associated baseline may be used for comparison with other baselines. However, changes associated with window 21 may not enter into changes associated with other windows, for example. Moreover, window 21 may be based on a preceding window, and the window it is based on may be changed over time as desired by a user.

[00242] Data from the olap cubes may be requested and received in a variety of ways. Understanding how requests and data gathering may interact may also help. Fig. 47 illustrates an embodiment of a process of responding to a request for data. Process 4700 includes requesting data, getting original data, gathering changes, and

displaying a response. Fig. 48 illustrates an embodiment of a process of storing a change in data. Process 4800 includes receiving a change, gathering surrounding data, and storing the change.

[00243] Thus, a change may be received at module 4810. Surrounding data is then collected (or comes with the change) at module 4820. Surrounding data to a change may be data such as a customer, user, window in which the change was made, product, region, time the change is applicable to (tomorrow, a forecast three months out, for example), and time the change is made, for example. The window may be designated by the user or may default to the window in which the system is to make changes at the time. The data of the change, both actual fact data and surrounding data, is stored at module 4830.

[00244] In the meantime, at module 4725, a request for data is received. At module 4740, original data related to the request is retrieved, such as from a main olap cube partition. At module 4750, changes to the data are gathered, such as at a change olap cube partition. The combined response is then displayed at module 4760. If the request includes an indication of a window to be used, data for that window is provided, along with data used as building blocks from prior windows. A default window may be imposed, including all current data, too.

[00245] The window specified may dictate that some changes are excluded, either because they are private (specific to a private window) or because the window is private and does not include all changes in a forecast. However, a request may also be set up to include a selection of windows, allowing a user to cut out changes made by the

executive suite or marketing group to determine if results are better that way, for example. Similarly, a response may provide a display of the changing information over time – indicating values at various windows, for example. Moreover, if the request of module 4725 is transmitted when changes are being stored at module 4830, the request may overlap the change storage, but the result may include the changes as long as sufficient time is allowed to collect changes.

[00246] Understanding how information may be applied in this system to various situations may provide further assistance. Fig. 49 illustrates an embodiment of a set of baselines and tracked data. System 4900 includes atomic data and various associated baselines. Data points 4910 are organized for convenience based on users who would use them. Thus, a ceo has an analyst target, a personal target and an internal target. A vp of sales has an official target, a personal target, a sales manager (internal sales group) target and a representative (external sales group) target. Similarly, a regional sales manager may have an official target and a personal target.

[00247] Various baselines are prepared. Sales baseline 4950 and company baseline 4930 are public baselines for use within the company for forecasts. Regional manager baseline 4960, vp of sales baseline 4940 and ceo baseline 4920 are personal baselines created to track alternative scenarios or goals which the various users do not wish to have public. Thus, the personal targets are associated with private baselines, and the various public targets are associated with public baselines.

[00248] However, these various numbers may each build off of a common history and common set of numbers which are aggregated to arrive at the various

expectations. Similarly, the baselines may build selectively off of other baselines – and baselines in general may be expected to depend from earlier baselines in terms of the data used and drawn from for calculations. Thus, each user may have a unique view of the data and still all look at (essentially) the same data. If baselines selectively build from earlier baselines, the information about which baselines are included will generally be accessible when such baseline information is used.

[00249] A system may be used to implement the various embodiments, allowing for real-time forecasting. Fig. 50 illustrates an embodiment of a medium and associated devices. System 5000 includes medium 5010, olap engine 5080 and cube 5070, database 5060, and user device 5090. Medium 5010 includes a database interface 5040, an olap interface 5030, a user interface 5050, and a control module 5020. Olap engine 5080 and olap cube 5070 provide the olap data structure, which is populated from the data in database 5060 using olap engine 5080. User device 5090 provides the user entry into the system, for queries and responses.

[00250] Thus, a query may go to user interface 5050, and then to control 5020. The query may then go to olap interface 5030, to get data from olap cube 5070, and then relay that data through control module 5020 and user interface 5050 to the user. Similarly, new data or change data may go from user interface 5050 to control module 5020 and then to database interface 5040. Database interface 5040 may then submit the data to database 5060, with the data then populated into a change partition of olap cube 5070. Eventually, the partitions of olap cube 5070 are unified, and the change data is thereby integrated.

[00251] From the foregoing, it will be appreciated that specific embodiments of the invention have been described herein for purposes of illustration, but that various modifications may be made without deviating from the spirit and scope of the invention. In some instances, reference has been made to characteristics likely to be present in various or some embodiments, but these characteristics are also not necessarily limiting on the spirit and scope of the invention. In the illustrations and description, structures have been provided which may be formed or assembled in other ways within the spirit and scope of the invention. Moreover, in general, features from one embodiment may be used with other embodiments mentioned in this document provided the features are not somehow mutually exclusive.

[00252] In particular, the separate modules of the various block diagrams represent functional modules of methods or apparatuses and are not necessarily indicative of physical or logical separations or of an order of operation inherent in the spirit and scope of the present invention. Similarly, methods have been illustrated and described as linear processes, but such methods may have operations reordered or implemented in parallel within the spirit and scope of the invention. Accordingly, the invention is not limited except as by the appended claims.

IN THE CLAIMS

We claim:

1. A system comprising:

a first client;

an analysis server incorporating an OLAP cube coupled to the first client; and

a first customer database of information coupled to the analysis server, the first customer database is to embody forecast data and to receive essentially real-time updates to the forecast data, the OLAP cube to analyze and update the information of the first customer database.

2. The system of claim 1, wherein:

the OLAP cube includes a main partition and a change partition, changes of the database incorporated into the change partition.

3. The system of claim 2, wherein:

new data of the database is entered with a change flag when such new data is submitted.

4. The system of claim 3, wherein:

The change partition is aggregated by the OLAP cube on a first predetermined time interval, the aggregating of the change partition does not include aggregating entries of the change partition into the main partition.

5. The system of claim 4, wherein:

The main partition is aggregated by the OLAP cube on a second predetermined time interval, the second predetermined time interval greater than the first predetermined time interval, the aggregation of the main partition includes incorporating the entries of the change partition.

6. The system of claim 1, wherein:

data is entered in the database with a timestamp and a public window designation associated therewith.

7. The system of claim 6, wherein:

baselines are maintained within the database, the baselines including data based on the window value associated with the data and window values designated for the baseline.

8. The system of claim 1, further comprising:

A second client coupled to the identification server and the analysis server; and

A second customer database of information coupled to the identification server and the analysis server, the second customer database dedicated to use by a second set of users of the system.

9. The system of claim 1, wherein:

The information is financial information.

10. The system of claim 1, further comprising:

means for providing data to the first client from a location remote from a location of the first client.

11. A method of maintaining information, comprising:

receiving a set of forecast data;

incorporating the forecast data into a database of the information through an OLAP cube;

extracting a baseline forecast from the database;

receiving updates to the database; and

propagating updates in a change partition of the OLAP cube.

12. The method of claim 11, further comprising:

flagging data updated in the database with a change flag.

13. The method of claim 12, further comprising:

aggregating changes from the change partition into a main partition of the OLAP cube;

and

unflagging the change flag of data updated in the database.

14. The method of claim 11, further comprising:

watching a set of watched data points of the information; and

notifying a user of changes in the set of watched data points responsive to the propagating and receiving updates.

15. The method of claim 11, wherein:

the method is executed by a processor in response to instructions, the instructions embodied in a machine-readable medium.

16. The method of claim 11, further comprising:

designating a window value when updates to the database are received, the window value associated with the updates;

recording a timestamp associated with updates to the database;

and

deriving a baseline from values of data in the database, the data selected according to a window value designated for the baseline.

17. The method of claim 16, wherein:

the baseline is derived from a previous baseline and updates received in the database since derivation of the previous baseline.

18. The method of claim 16, wherein:

the baseline is derived responsive to a user request for a private baseline incorporating a subset of the data of the database, the subset determined based on window value associated with the data.

19. The method of claim 16, wherein:

the baseline is derived from a subset of a set of previous baselines and updates received in the database since derivation of the previous baseline having window values of a selected subset of available window values.

20. The method of claim 11, wherein:

The forecast data is financial information.