

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
11 December 2008 (11.12.2008)

PCT

(10) International Publication Number
WO 2008/150533 A2

- (51) International Patent Classification:
G06Q 40/00 (2006.01)
- (21) International Application Number:
PCT/US2008/007025
- (22) International Filing Date: 2 June 2008 (02.06.2008)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
60/941,334 1 June 2007 (01.06.2007) US
- (71) Applicant (for all designated States except US): FTEN, INC. [US/US]; 800 Third Avenue, Twenty- Third Floor, New York, NY 10022 (US).
- (71) Applicant and
- (72) Inventor (for all designated States except US): KITTELSEN, Douglas, G. [US/US]; 29824 Paintbrush Drive, Evergreen, CO 80439 (US).
- (74) Agents: HSU, Lin, C. et al.; Holme Roberts & Owen LLP, 299 South Main Street, Suite 1800, Salt Lake City, UT 84111 (US).

- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

[Continued on next page]

(54) Title: METHOD AND SYSTEM FOR MONITORING MARKET DATA TO IDENTIFY USER DEFINED MARKET CONDITIONS

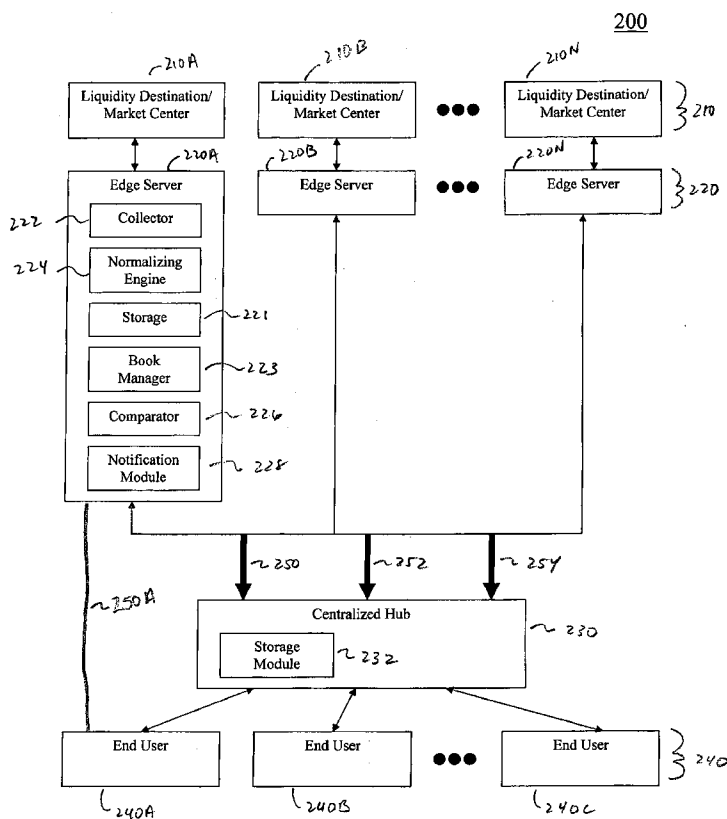


FIG. 2

(57) Abstract: A method and system for monitoring market data. The method includes collecting real time data that is related to conditions of a trading market. Collection occurs at an edge server associated with a liquidity destination trading at least one financial article of trade. In addition, the real time data that is collected is also normalized into a standard form. A user defined criteria is received from a centralized hub. The user defined criteria defines a particular event in the condition. It is then determined when a condition in the trading market matches the event. A response is generated providing notification of the occurrence of the event. The response is sent to the centralized hub for distribution to a user associated with the user defined criteria.

WO 2008/150533 A2



Published:

- *without international search report and to be republished upon receipt of that report*

METHOD AND SYSTEM FOR MONITORING MARKET DATA TO IDENTIFY USER DEFINED MARKET CONDITIONS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] The present application claims priority to and the benefit of Provisional Application No. 60/941,334 to Kittelsen, entitled "EdgeXpress™ Market Data System," filed on June 1, 2007, which is herein incorporated by reference in its entirety.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] The present invention relates to a system for consolidating and analyzing intraday issues related to securities transactions that have been submitted to liquidity destinations, and more specifically to a system and method for monitoring market data to identify user defined market conditions using a distributed system of edge servers located at each of a plurality of liquidity destinations and a centralized hub for interfacing end users with the edge servers.

2. The Relevant Technology

[0003] In today's financial securities industry, advances in technology and high-speed, high volume computerized algorithmic trading strategies have combined to make optimized speed critical to success. The markets began as a totally manual process where trade volumes were very small and the time to find a match was measured in minutes. With the advent of electronic alternative trading networks in the 1990's, much more volume was handled by computers without human intervention, which caused the relevant industry measure to move from minutes to seconds, from seconds to milliseconds, and most recently from milliseconds to microseconds.

[0004] Not only have the expectations of market processing time changed radically, but the number of total orders that are processed has grown exponentially. This growth in total order volume is driven by high-speed, high volume computerized algorithmic trading models that literally flood the markets by placing thousands of orders and cancels for orders per second as a means to exploit momentary imbalances within and/or between various liquidity destinations/market centers to attempt the execution of "lightning quick" purchases.

[0005] Tremendous time and resources are invested by industry participants to minimize the length of time it takes to send transaction information to and from liquidity destinations/market centers. As a result, the outer most limits of performance improvements available from using more powerful computer equipment and telecommunications capabilities are constantly being stretched. In addition, many of the high-speed, high-volume trading strategies include high cancellation rates if all the desired elements of a transaction are not present. This results in as many as one-thousand potential transactions being sent to liquidity destinations/market centers for each transaction that is executed. Moreover, each order placed and cancelled generates additional market data that needs to be analyzed, which in turn may generate more market data, thus creating a feedback loop. As a result, this high volume of messages being sent to and from various liquidity destinations/market centers puts tremendous strain on telecommunications capabilities and creates "queuing" delays as messages are forced to wait until previously transmitted messaging traffic is processed.

[0006] Minimizing the time necessary to send transaction information to liquidity destinations/market centers is only part of the challenge. The market factors outlined above also result in a tremendous volume of message traffic emanating from the liquidity destinations/market centers related to the then-current state of the market for each of the

numerous stock symbols and associated transactions ("market data"), all with widely varying degrees of importance and relevance.

[0007] For instance, PRIOR ART FIG. 1 illustrates how market data is aggregated together indiscriminately and sent down common communications lines in a typical data processing system. This aggregation causes significant delays and latency in delivering the market data necessary for market participants to make trading decisions in a timely manner. In particular, in the typical model illustrated in FIG. 1, the servers receiving the market data cannot keep up with the increasing flow of information. As a result, the incoming data is queued for processing, leading to additional latency. Also, once the queue has reached maximum capacity, market data can be dropped, which causes the loss of data and therefore the true understanding of the status of a symbol.

[0008] As shown in FIG. 1, the current industry approach for processing market data is to establish direct connectivity with the different exchanges and bring all the information that the exchanges publish to a group of servers that process the data needed by users or computerized algorithmic trading programs. The key drawbacks with this approach is the extremely high cost of transporting and processing such a huge amount of data, and the ability to do so in a timely fashion. Transmitting this large amount of aggregated, unanalyzed, non-normalized data introduces latency and delays in transmission times. In addition, only after this data is received in a common repository can it then be analyzed, which adds additional time, before it can be sent to end users according to their specific interests.

[0009] Further, as illustrated in PRIOR ART FIG. 1, existing market data systems can be analogized to sending a large number of multi-colored glass beads of varying sizes and colors emanating from multiple sources down a common aggregation funnel which is too small to

receive all of the beads at one time so the beads queue up. Eventually all the glass beads will pass through the funnel and arrive at a common repository, at which point they can be analyzed and separated by color and size. If a user was interested in only seeing green glass beads of a specific size, they would have to wait for all of the different glass beads to pass through the tunnel and then wait for the proper size green glass beads to be separated from the rest, thereby increasing time for processing the market data.

SUMMARY OF THE INVENTION

[00010] A method and system for monitoring market data. The method includes collecting real time data that is related to conditions of a trading market. Collection occurs at an edge server associated with a liquidity destination trading at least one financial article of trade. The real time information is stored at the edge server. In addition, the real time data that is collected is also normalized into a standard form at the edge server. A user defined criteria is received from a centralized hub and is delivered and subsequently processed at the edge server. The user defined criteria defines a particular event in the condition. It is then determined when a condition in the trading market matches the event. A response is generated providing notification of the occurrence of the event. The response is sent to the centralized hub for distribution to a user associated with the user defined criteria.

BRIEF DESCRIPTION OF THE DRAWINGS

[00011] Exemplary embodiments are illustrated in referenced figures of the drawings which illustrate what is regarded as the preferred embodiments presently contemplated. It is intended that the embodiments and figures disclosed herein are to be considered illustrative rather than limiting.

[00012] PRIOR ART FIG. 1 is a diagram of a typical market data system that introduces latency when processing market data.

[00013] FIG. 2 is an illustration of a market data system configured such that edge servers located at the liquidity destinations are able to maintain real-time market books by symbol and to identify when user defined conditions occur in the market, in accordance with one embodiment of the present invention.

[00014] FIG. 3 is a flow diagram illustrating a method for monitoring market data from the viewpoint of an edge server at a liquidity destination, in accordance with one embodiment of the present invention.

[00015] FIG. 4 is a flow diagram illustrating a method for monitoring market data from an overall viewpoint of a system capable of implementing the method, in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

[00016] Reference will now be made in detail to the preferred embodiments of the present invention, a method and system for the monitoring of market data to identify user defined market conditions using a distributed system of edge servers located at each of a plurality of liquidity destinations and a centralized hub for interfacing end users with the edge servers. While the invention will be described in conjunction with the preferred embodiments, it will be understood that they are not intended to limit the invention to these embodiments. On the contrary, the invention is intended to cover alternatives, modifications and equivalents which may be included within the spirit and scope of the invention as defined by the appended claims.

[00017] Accordingly embodiments of the present invention provide for the collection and normalization of market data at their source. In addition, other embodiments of the present invention are capable of distinguishing from a wide variety of events reflected in the market data. As such, certain user defined events are capable of being identified, and notification of the occurrence of the event is sent to the requesting user. Also, still other embodiments provide the added benefit of analyzing market data without performing any preliminary analysis of the relevance of different information to different end users.

Notation and Nomenclature

[00018] Embodiments of the present invention can be implemented on a software program for processing data through a computer system. The computer system can be a personal computer, notebook computer, server computer, mainframe, networked computer (e.g., router), handheld computer, personal digital assistant, workstation, and the like. This program or its corresponding hardware implementation is operable for enabling the monitoring of market data. In one embodiment, the computer system includes a processor

coupled to a bus and memory storage coupled to the bus. The memory storage can be volatile or non-volatile and can include removable storage media. The computer can also include a display, provision for data input and output, etc.

[00019] Some portion of the detailed descriptions that follow are presented in terms of procedures, steps, logic block, processing, and other symbolic representations of operations on data bits that can be performed on computer memory. These 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. A procedure, computer executed step, logic block, process, etc. is here, and generally, conceived to be a self-consistent sequence of operations or instructions 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 in a computer system. 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.

[00020] 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 discussions, it is appreciated that throughout the present invention, discussions utilizing terms such as “defining,” “receiving,” “determining,” “comparing,” or the like refer to the actions and processes of a computer system, or similar electronic computing device, including an embedded system, that manipulates and transfers 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.

System for Monitoring Market Data

[00021] FIG. 2 is a data flow diagram illustrating the flow of information in a data processing system 200 capable of monitoring market data from a plurality of liquidity destinations/market centers 210. The subject of this application is a financial article of trade data processing system 200 including various functions, which may be implemented to perform consolidation and analysis of intraday issues related to submitted securities, commodities, options, futures transactions. The system 200 adapts the submitted securities, commodities, options, futures, or other financial articles of trade transaction to involve at least one of disparate systems, multiple parties and multiple liquidity destinations 210.

[00022] Reference to securities or security transactions within this disclosure should be interpreted as transactions involving securities, commodities, option or futures. The term other financial articles of trade transaction refers to any other article traded at the liquidity destinations other than securities, commodities, options, futures. It should be appreciated that different configurations can be used to consolidate and analyze the data to achieve a given result. Although particular combinations are disclosed, variations on those combinations can be used to achieve the same consolidation and analysis in the financial article of trade data processing system 200.

[00023] As such, embodiments of the present invention enables users to take action promptly on an intraday basis to take advantage of real-time market conditions. In particular, the financial article of trade data processing system 200 is capable of identifying particular conditions within the market that are of interest to users of the processing system 200. As a result, once those market conditions are identified and notification of such is presented, the

users can react accordingly to take advantage of those market conditions. For instance, a Strat 1 coordinated time stamp is added to each message received from the liquidity destination. As such, a comparison can be made to published time stamps from the liquidity destination to determine if there is a drift. If there is a difference in time stamps this is an indicator that the exchange is under load and not processing the market flow efficiently. This may create an opportunity to trade. In addition, this information is useful to ensure that there are no communication problems between the edge and a centralized hub.

[00024] As shown in FIG. 2, embodiments of the present invention tackle the problem of handling the massive amounts of data by reversing the approach used to gather and transmit the data to be acted upon. In particular, rather than simply shipping all possible market data from the plurality of liquidity destinations 210 to a central location, dedicated computer processing power is placed as close to each liquidity destinations/market centers as possible. These servers are referred to as edge servers for purposes of the present Application. For instance, each of the plurality of liquidity destinations 210 in the data processing system 200 is associated with one of a plurality of edge server 220. For instance, liquidity destination 210A is associated with edge server 220A, liquidity destination 210B is associated with edge server 220B, and liquidity destination 210N is associated with edge server 220N.

[00025] While FIG. 2 shows one edge server associated with a corresponding liquidity destination, it is to be understood that one or more edge servers may support a single liquidity destination. For instance, different edge servers may be responsible for different ranges of symbols to distribute the load coming from the corresponding liquidity destination. Additionally, in other embodiments, a liquidity destination also represents a “dark pool” that is not an exchange, or a crossing engine (e.g., for internal crossing or matching inside a large

clearing company, or any place that trade matching can occur and that publishes market data, or accepts indication of interest (IOI) messages.

[00026] As such, by locating dedicated computer processing resources (e.g., edge servers 220A-N) at or communicatively near the source of market data (e.g., liquidity destination 210A-N), it is possible and affordable to have huge amounts of bandwidth to receive all the data generated by the exchange to the edge server for initial processing. The initial processing performed at each of the plurality of edge servers 210 will filter and normalize the data for this exchange as well as maintain a real-time market books by symbol (e.g., stock market symbol or name). Since this server is dedicated solely to a specific exchange, each available CPU cycle can be used to process and prepare the market data. In a case with multiple CPUs supporting a specific exchange, many more cycles are available to process and prepare the market data.

[00027] As an example, the components and features of edge server 220A is representative of the components and features of each of the plurality of edge servers 210. For instance, edge server 220A includes a collector 222 for collecting real time data at a corresponding liquidity destination. In this case, edge server 220A collects real time data from liquidity destination 210A that is trading at least one financial article of trade. In particular, edge server 220A receives raw market data from the liquidity destination 210A. For instance, raw data includes submitted and executed securities transactions particular to liquidity destination 210A.

[00028] In addition, edge server 220A includes a normalizing engine 224, which converts the real time data that is collected into a standard format. Normalization of the market data allows for ready analysis with market data from other liquidity destinations in the plurality of liquidity destinations 210.

[00029] Edge server 220A also includes a storage module 221. The normalized information is stored in the storage module 221 for purposes maintaining a history of the market conditions reflected at the liquidity destination 210A. In addition, other information based on the normalized information is also stored in storage module 221, as will be described below.

[00030] For instance, book manager 223 is able to process the normalized market data and maintain a Level 2 book (depth of book) information for every symbol traded on liquidity destination 210A. Level 2 book information includes the number of shares offered at varying price points that is different from the current bid and asking price for a particular symbol. In addition, each symbol may be associated with one or more standard industrial classification (SIC) codes. As such, it is possible to monitor one or more SIC codes, instead of the symbols. This allows for the monitoring of a family of related symbols based on the industries they in which they participate.

[00031] In addition, from the Level 2 book information, the edge server 220A is able to determine the current Level 1 book information (e.g., current bid and ask pricing) for each of the symbols traded on the liquidity destination 210a. This information is stored and also delivered to the centralized hub 230 for distribution to any interested users in the plurality of users 240.

[00032] As the edge server 2120A receives and normalizes the data for liquidity destination 210A, it also ships the relevant data to the centralized hub 230. As shown in FIG. 2, the hub 230 is coupled to each of the plurality of edge server 220, and in particular includes real-time communicative coupling to each edge server.

[00033] For instance, three independent data paths are enabled for each edge server, such as edge server 220A. Data path 250 is dedicated to the communication of real time, one-to-one message traffic from the edge server 220A to centralized hub 230. That is, as messages are generated by liquidity destination 210A and received by the edge server 220A, those messages are normalized and sent directly to the centralized hub 230 via data path 250, for distribution to interested users in the plurality of end users 240.

[00034] In another embodiment, the data path 250 includes a multicast channel 250A that is accessible to authorized users. FIG. 2 shows one such multicast channel 250A coupled between edge server 220A and end user 240A. Additional multicast channels coupling various edge servers and users are not shown in FIG. 2. As such, the normalized data communicated over data path 250 is passed from an edge server to the end user, without necessarily having to route through the centralized hub 230.

[00035] In addition, data path 252 is dedicated to signal management. In particular, attributes representing certain market conditions as defined by an end user are delivered to the centralized hub 230. For instance, end user 240A is interested in determining if and when a particular market condition exists for one or more symbols or SIC codes traded on the plurality of liquidity destinations 210. This desire is expressed in a request for notification of the condition that is delivered to the centralized hub 230 from a requesting user. This market condition can be represented by one or more attributes of a proprietary protocol that is recognized by system 200, and in particular, by hub 230 and each of the plurality of edge server 220. These attributes can focus on any aspect of the market data being processed and accumulated at the edge server 220A, and can include a specific symbol, a group of symbols, price changes of symbol(s), time drift that senses a market is overloaded and responding slowly, etc. In addition, the attributes can include and describe the available liquidity on the

market, as implied in Level 2 information, as part of the data that is used for defining a market condition. As a result, the market condition defines an event that is further defined by one or more attributes, as outlined in the request delivered to the hub 230 from the end user 240A. When events in the market data match the request, the user will receive a response and can make a trading decision based on the knowledge that the event has occurred. In one embodiment, the trade decision is automatically processed based on predefined instructions.

[00036] More particularly, the request is delivered from the centralized hub 230 to the corresponding edge server supporting the liquidity destination that is trading on the one or more symbols, such as edge server 220A over path 252 via data path 252. Moreover, the hub 230 persists the monitoring request and passes the request to the particular venues named in the request. For instance, in the case of edge server 220A, it creates the necessary mapping rules to evaluate the market data received that that point forward to see if it matches the pattern requested in the request. That is, comparator 226 performs the functions necessary to identify when events match the requested market conditions.

[00037] In addition, centralized hub 230 also maintains a cache of all open requests so it can query the current active signals for a requesting customer. That is, hub 230 is able to determine the status of a particular request as monitored by each of the plurality of edge server 220.

[00038] Because the centralized hub 230 acts as an interface between the plurality of end users 240 and the plurality of edge servers 220, a user can add, delete, or update a monitoring request at any time in the trading day. As such, requests can be updated to reflect current market conditions.

[00039] Furthermore, data path 254 is dedicated to the communication of results from the request. That is, once an event is recognized by the comparator 226 of edge server 220A, notification of the occurrence of the event is presented to the centralized hub via data path 254. For instance, notification module 228 generates the notification of the occurrence of the event, and also delivers the notification to the centralized hub 230 via data path 254 for distribution to interested end users. In addition, as signal responses (e.g., notifications) are received from edge server 220A, they are stored in a local cache 232. Because the data path 254 is dedicated to communicating results, receiving systems can give signal responses the highest level of priority.

[00040] The centralized hub 230 acts as the coordinator that dynamically determines what data must be shipped between the plurality of edge servers 220 and the centralized hub 230, based on the specific needs of each of the plurality of end users 240. For instance, the centralized hub 230 receives requests for particular information from each of its consumers and then immediately formulates how to gather the appropriate data to meet that specific user's request. Since the plurality of edge servers 220 have already normalized the data and maintained their own market book for each symbol, the centralized hub 230 can simply request the specific information that it needs for the specific user, in one embodiment. That data and subsequent updates to that data will automatically flow from all relevant edge servers to the centralized hub 230 for final merging, consolidation, and delivery to interested users on an ongoing basis.

[00041] For instance, in one embodiment, the centralized hub 230 acts as a comparator when determining if a cross-venue condition exists. In this case, the comparator at any edge server would not be able to determine the existence of the cross-venue condition for lack of information; however, the centralized hub is in a position to recognize the cross-venue

condition. Specifically, a monitoring request is generated by a user. The monitoring request enables the monitoring of events occurring at more than one venue (liquidity destination). As previously described, the centralized hub 230 is able to determine which information is necessary from each of the venues in question. As such, the centralized hub 230 is able to request the necessary information from each of the necessary edge servers. Moreover, the centralized hub 230 aggregates individual signals from the corresponding edge servers. Further analysis is performed by the centralized hub 230 to determine if the individual signals in combination indicate that the cross-venue condition exists. Only when the cross-venue condition exists will the centralized hub 230 generate the notification signal for delivery to the appropriate user associated with the request.

[00042] In one particular embodiment, a user can register with the centralized hub 230 to monitor a list of symbols being traded in corresponding liquidity destinations. As a result, data processing system 200 will send snapshots of a corresponding symbol that describes the latest changes of the symbol over a given time period. These snapshots are used to determine a current state of a corresponding symbol. As such, the user is able to determine a proper course of action with regards to the market conditions for the corresponding symbol.

[00043] As a representative example, a symbol being traded on liquidity destination 210A is the subject of snapshot. As such, edge server 220A will take snapshots of the current state of the corresponding symbol every x microseconds (where x is configurable). These snapshots provide updates to the previously known state of the corresponding symbol, and is sent to the centralized hub 230. In particular, the snapshots reflect the net changes that have occurred since the last snapshot. If the book has not changed since the last snapshot - no new snapshot will occur. If the liquidity destination 210A is very busy - the impact off the snapshot approach is most beneficial. For instance, if in the snapshot window liquidity

destination 220A generated one-thousand messages, the edge server 220A would generate and send a single snapshot message, that includes the net effect of those one-thousand messages, to the centralized hub 230, for consolidation and updating of the corresponding market wide book.

[00044] As described, the centralized hub 230 will receive the snapshot from edge servers 220A-N supporting the plurality of liquidity destinations 210 and combine them into a consolidated book for each symbol. In addition, the centralized hub 230 also supports requests from a user (e.g., one of users 240A-N) asking to register interest in a symbol. If a user has registered interest in a name, the current book state is generated for the symbol, based on previously received snapshot information, and delivered to the user from the centralized hub 230. This allows a user to come into the markets mid-day and immediately get the current state of any name and receive all relevant snapshots until they deregister.

[00045] The result of this invention is the fastest possible delivery of relevant market data to consumers by sending only what is necessary, preprocessing the data as it appears at the exchange in order to streamline the merging of data from the many exchanges and reducing the total amount of bandwidth that is necessary. In addition, as the volume of market data increases on the liquidity destinations 210, the time and data processing saving benefits provided by the plurality of edge servers 220 will be self evident.

Method for Monitoring Marketing Data

[00046] FIG. 3 is a flow diagram 300 illustrating a method for monitoring market data from the viewpoint of an edge server associated with a liquidity destination, in accordance with one embodiment of the present invention. The method of 300 is implemented by the data processing system 200 of FIG. 2 in one embodiment, and more particularly, by each of

the plurality of edge servers 220. That is, the functionality of an edge server is described in flow diagram 300.

[00047] At 310, real time data is collected at the edge server that is associated with a corresponding liquidity destination. The real time data is related to market conditions at the corresponding liquidity destination trading at least one financial article of trade. For instance, the real time data collected relates to Level 2 depth of book information for symbols traded on the corresponding liquidity destination. As described previously, each liquidity destination is associated with an edge server for purposes of relieving the data flow presented to a centralized hub that interfaces with a plurality of end users. In one embodiment, collector 222 of edge server 220A in FIG. 2 implements the functionality of 310 to collect real-time data.

[00048] At 320, the real time data is normalized at the corresponding edge server. In particular, the real-time data is normalized into a standard form. That is, the real-time data is translated into an appropriate normalized format that is suitable for analysis and comparison to other data collected by other edge servers associated with other liquidity destinations. In this manner, a distributed system for the collection of market data is achieved, such that all market data for a particular symbol can be collected on various liquidity destinations and normalized for consolidation and analysis. In one embodiment, the normalizing engine 224 of FIG. 2 implements the functionality of 320.

[00049] At 330, the real time data that is normalized is stored at the corresponding edge server associated with the liquidity destination. In this manner, the history of the market conditions for a liquidity market is accessible at the corresponding edge server in case the information is lost at the centralized hub. In addition, analysis on the real time data that is normalized can be performed in order to recognize user defined market events, and for

purposes of sending snapshots of market conditions over a period of time. In one embodiment, the storage module 221 of FIG. 2 implements the functionality of 330.

[00050] At 340, user defined criteria is received from a centralized hub that interfaces with a plurality of end users. The user defined criteria is received at the corresponding edge server associated with the liquidity destination. The user defined criteria defines a particular event as described by market conditions at the corresponding liquidity destination. For instance, the event may involve information related to state of book for one or more symbols traded on the liquidity destination, or be based on particular parts of a book (e.g., the top three price points). As such, embodiments of the present invention are able to implement a system for identifying when a predefined event occurs, wherein the identification is performed at the edge server associated with the corresponding liquidity destination. Comparator 226 is capable of performing the operation at 340.

[00051] At 350, a condition in the trading market associated with the liquidity destination is compared to the event. For instance, comparator 226 performs the operation at 350, in one embodiment. In this manner, instead of only sending normalized data to the centralized hub, or to the end user for further analysis, the edge server is able to identify particular events that are important to an end user.

[00052] At 360, the edge server generates a response that provides notification of the occurrence of the event. More particularly, a notification signal that is recognizable by the user requesting the notification is generated. In one case, the notification signal is provided by the user in the original request. As such, the user is able to define the format of the signal response, or notification signal, and the underlying name of the signal to match their particular processing needs. Upon receipt of the notification signal, the user is able to

immediately recognize that the event of interest, as described by market conditions, has occurred. The user is able to react accordingly in a time sensitive fashion.

[00053] At 370, the response is sent to the centralized hub for distribution to the interested user making the original request. The centralized hub acts as the interface between the end users and the edge servers supporting the plurality of liquidity destinations. The response includes the necessary information for the centralized hub to correctly direct the information to the appropriate user.

[00054] FIG. 4 is a flow diagram 400 illustrating steps in a method for providing snapshots of market conditions, in accordance with one embodiment of the present invention. The operations provided in flow diagram 400 is implemented by system 200 of FIG. 2, in one embodiment.

[00055] Operations outlined in 410, 420, and 430 are analogous to operations 310, 320, and 330, respectively, as previously described in relation to FIG. 3. The operations in flow diagram 400 are applied to a plurality of edge servers supporting a plurality of liquidity destinations, instead of a single edge server supporting a corresponding liquidity destination, as described in FIG. 3. As such, the description in relation to 310, 320, and 330 is applicable to operations 410, 420, and 430 for each of the edge servers. In particular, each of the plurality of edge servers collect real-time data that is related to market conditions. In addition, the real-time data that is collected is further normalized into a standard format for analysis and comparison such that market data collected across a distributed system of edge servers can be compared readily in a common format. In addition, the real-time data that is normalized is stored at corresponding edge servers.

[00056] At 440, a plurality of snapshots of symbols traded on the plurality of liquidity destinations is generated at each of the plurality of edge servers. The snapshot comprises a net effect of trading activity on a corresponding symbol for a given time period. In addition, the plurality of snapshots is communicated to the centralized hub that interfaces with a plurality of end users. In one embodiment, if the net effect is approximately zero, then the snapshot is not communicated to the centralized hub. As such, in conjunction with 440, the plurality of snapshots is also received at the centralized hub.

[00057] In this manner, the centralized hub is able to consolidate all the snapshots received from the plurality of edge servers supporting a plurality of liquidity destinations to provide a consolidated view of the symbols traded on the liquidity destinations. That is, at 450, books of the symbols traded on the liquidity destinations are consolidated using the plurality of snapshots. For instance, Level 1 and Level 2 book information from all the liquidity destinations are consolidated at the centralized hub. More particularly, each snapshot is combined with a previous state of the corresponding book of a corresponding symbol. In this, manner the corresponding book is updated to obtain a current state. The state of the book is available for distribution to interested users.

[00058] Accordingly embodiments of the present invention provide for the collection and normalization of market data at their source. In addition, other embodiments of the present invention are capable of distinguishing from a wide variety of events reflected in the market data. As such, certain user defined events are capable of being identified, and notification of the occurrence of the event is sent to the requesting user.

[00059] While the methods of embodiments illustrated in flow charts 3 and 4 show specific sequences and quantity of operations, the present invention is suitable to alternative embodiments. For example, not all the operations provided for in the methods presented

above are required for the present invention. Furthermore, additional operations can be added to the operations presented in the present embodiments. Likewise the sequences of operations can be modified depending upon the application.

[00060] A method and system for the monitoring of market data to identify user defined market conditions using a distributed system of edge servers located at each of a plurality of liquidity destinations and a centralized hub for interfacing end users with the edge servers, is thus described. While the invention has been illustrated and described by means of specific embodiments, it is to be understood that numerous changes and modifications may be made therein without departing from the spirit and scope of the invention as defined in the appended claims and equivalents thereof. Furthermore, while the present invention has been described in particular embodiments, it should be appreciated that the present invention should not be construed as limited by such embodiments, but rather construed according to the below claims.

CLAIMS

We claim:

1. A method for monitoring market data, comprising:
 - collecting real time data related to conditions of a trading market at an edge server associated with a liquidity destination trading at least one financial article of trade;
 - normalizing said real time data that is collected into a standard form at said edge server;
 - storing said real time data that is normalized at said edge server;
 - receiving from a centralized hub user defined criteria defining a particular event in said conditions at said edge server;
 - determining if a condition in said trading market matches said event;
 - generating a response providing notification of the occurrence of said event; and
 - sending said response to said centralized hub for distribution to a user associated with said user defined criteria.

2. The method of Claim 1, wherein said collecting real time data comprises:
 - collecting Level 2 depth of book information as said real time data for symbols traded on said liquidity destination.

3. The method of Claim 2, further comprising:
 - analyzing said Level 2 book information to determine Level 1 top of book information for symbols traded on said liquidity destination; and

sending said Level 1 top of book information to said centralized hub for distribution to interested users.

4. The method of Claim 1, wherein said generating a response comprises:
generating a notification signal.
5. The method of Claim 4, further comprising:
receiving said notification signal from said user, wherein said user defines said notification signal to indicate said event has occurred.
6. The method of Claim 1, further comprising:
sending said real time data to a centralized hub in a one-to-one message ratio.
7. The method of Claim 1, further comprising:
for a given time period, determining a net effect of trading activity on a symbol traded on said liquidity destination; and
sending a snapshot of said symbol including said net effect to said centralized hub for consolidation and updating, wherein said snapshot is combined with a previous state of a book of said symbol at said hub to update said state.
8. The method of Claim 6, wherein said sending a snapshot further comprises:
declining to send said snapshot if said net effect is zero.
9. The method of Claim 1, wherein said event is related to a particular symbol.

10. The method of Claim 1, wherein said event is related to two or more symbols.
11. The method of Claim 1, wherein said determining if a condition comprises:
comparing time stamps provided by said liquidity destination and as reported to
determine if the market is falling behind.
12. The method of Claim 1, further comprising:
allowing said user to define said user defined criteria in a request;
accepting said request at said centralized hub;
updating said request; and
distributing said request to said edge server.
13. The method of Claim 1, wherein said event is related to an SIC code.
14. A method for monitoring market data, comprising:
at each of a plurality of edge servers associated with a plurality of liquidity
destinations trading at least one financial article of trade, collecting real time data related to
market conditions;
at each of said plurality of edge servers, normalizing said real time data that is
collected to a standard form;
at each of said plurality of edge servers, storing said real time data that is normalized;
at a centralized hub, receiving a request from a user for monitoring an event in said
market conditions, wherein said request includes user defined criteria that defines said event;
sending said user defined criteria to pertinent edge servers associated with
corresponding liquidity destinations as defined in said request;

receiving a response to said request at said centralized hub indicating an occurrence of said event; and

sending said response to said user.

15. The method of Claim 14, wherein said receiving a response comprises:
receiving a notification signal as said response, wherein said notification signal is defined by said user and previously sent to said pertinent edge servers.

16. The method of Claim 14, further comprising:
receiving said real time data that is normalized from each of said plurality of edge servers at said centralized hub in an approximate one-to-one message ratio for distribution to interested users.

17. The method of Claim 14, further comprising:
storing said request at said centralized hub; and
querying said pertinent edge servers for updates on said request.

18. The method of Claim 14, further comprising:
receiving a plurality of snapshots of symbols traded on said plurality of liquidity destinations, wherein a snapshot comprises a net effect of trading activity on a corresponding symbol for a given time period; and

consolidating books of said symbols traded on said plurality of liquidity destinations using said plurality of snapshots by updating said books, wherein said snapshot is combined with a previous state of a corresponding book of said corresponding symbol for updating a current state said corresponding book for distribution to interested users.

19. The method of Claim 18, wherein said consolidating books comprises:
consolidating Level 1 and Level 2 book information for said symbols traded on said plurality of liquidity destinations.

20. The method of Claim 14, further comprising:
receiving Level 1 book information from said plurality of edge servers for symbols traded on said plurality of liquidity destinations for distribution to interested users.

21. A system for monitoring market data, comprising:
a plurality of liquidity destinations trading at least one financial article of trade;
a plurality of edge servers associated with said plurality of liquidity destinations,
wherein at least one of said plurality of edge servers comprises:
a collector for collecting real time data related to market conditions;
a normalizing engine for normalizing said real time data that is collected to a standard form;
a comparator for monitoring an event in said market conditions as requested by a user;
a notification module for generating a notification signal providing notification of the occurrence of said event, and sending said notification signal to a centralized hub; and
said centralized hub for receiving a request from said user for monitoring said event,
wherein said request includes user defined criteria that defines said event, wherein said centralized hub is configured to send said request to pertinent liquidity destinations as defined

in said request, and wherein said centralized hub is configured to receive said notification signal for distribution to said user.

22. The system of Claim 21, wherein said centralized hub is configured to query edge servers associated with said pertinent liquidity destinations for updates on said request.

23. The system of Claim 21, wherein said centralized hub comprises a consolidator configured to receive a plurality of snapshots of a plurality of symbols traded on said plurality of liquidity destinations, wherein a snapshot comprises a net effect of trading activity on a corresponding symbol for a given time period, and configured to consolidate a books of said plurality of symbols traded on said plurality of liquidity destinations using said plurality of snapshots by updating said books, wherein said snapshot is combined with a previous state of a corresponding book of said corresponding symbol for updating a current state said corresponding book for distribution to interested users.

24. The system of Claim 21, wherein said centralized hub is configured to receive Level 1 book information from edge servers associated with said plurality of liquidity destinations for symbols traded on said plurality of liquidity destinations for distribution to interested users.

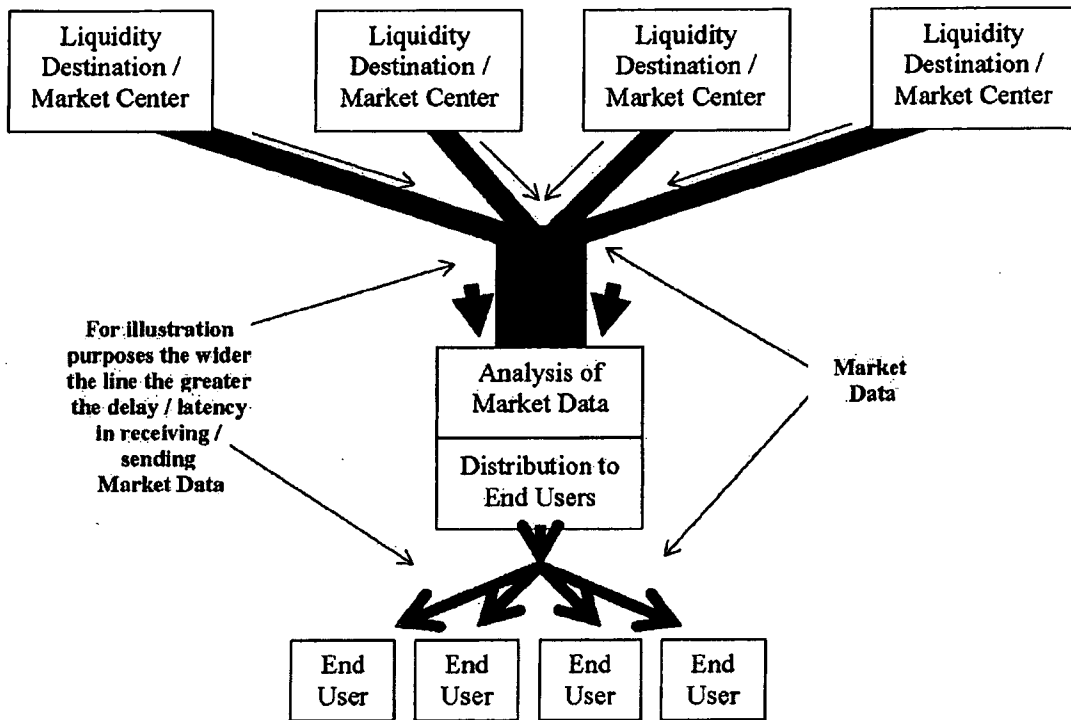


FIG. 1
PRIOR ART




FIG. 2

300

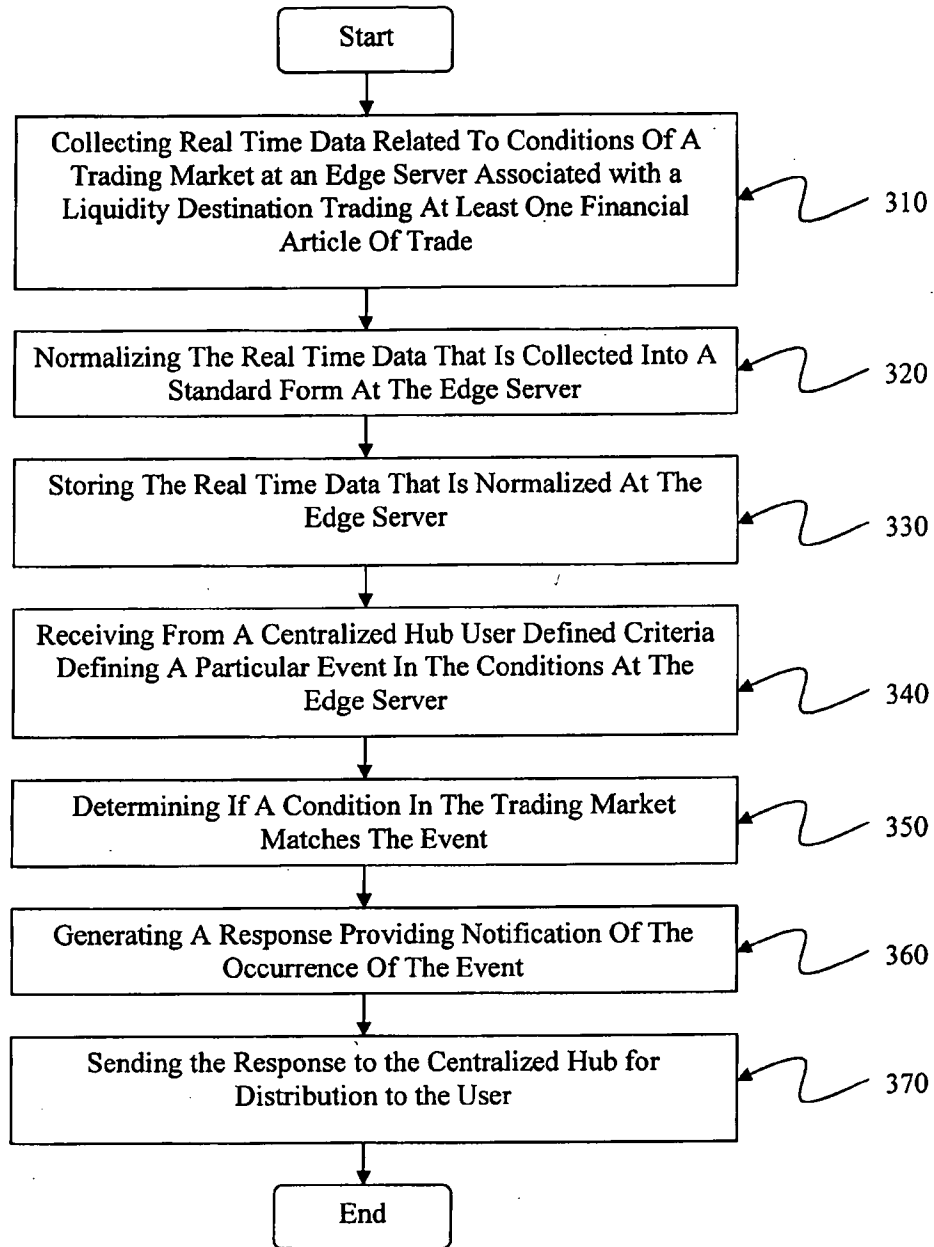


FIG. 3

400

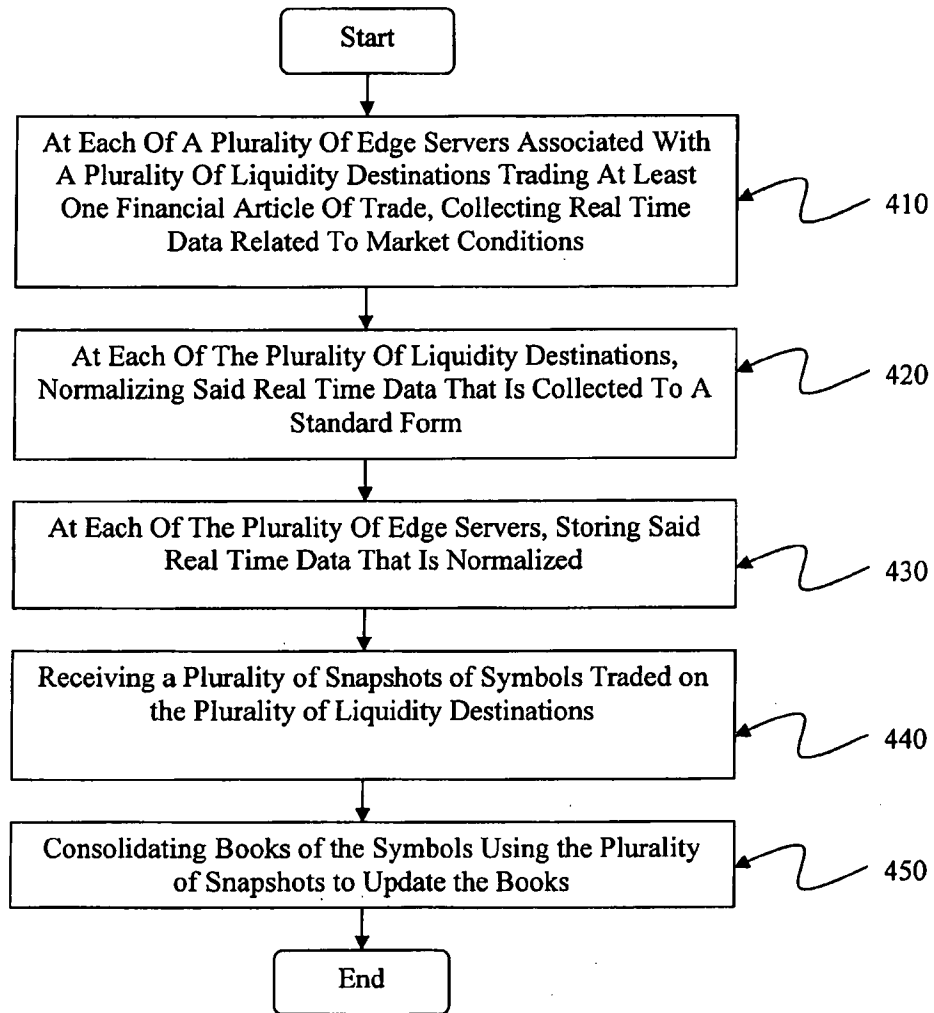


FIG. 4