

receiver 228 generally include information that allows the network controller 214 to track, among other things, a subscriber's program access history, as described below. The control receiver 228 can store the status reports internally in a local storage or memory device and transfer them to the network controller CPU 224. Typically, the control receiver 228 is interfaced with the network controller CPU 224 using standard RS-232 or RS-422 links, an IEEE-488 bus or the like.

10 In the preferred embodiment, the network controller CPU 224 scans the control receiver 228 at a predetermined rate (e.g., once every few seconds) to initiate the status report transfer. Upon transfer, the network controller CPU 224 adds the data and control information in the status reports to the network control databases 226 by: checking for changes in previously received status information, processing the new information and updating the corresponding parameters in the network control databases 226. The network controller 214 processes the information stored in its databases with any program control information relayed through the signal processor 209 from the delivery system's operations center 202. This processing capability allows the network controller 214 to modify prior control signals and create new ones. The network controller 214 transfers both modified and unmodified control signals, along with any local combined program signals 205, to the signal processor 209 to be

combined with others program signals 205 for distribution over the cable system 200.

3. Modifying the Program Control Information Signal

5

Tables A-C, described below, provide an example of some information that can be sent in the program control information signal to the set top terminals 220. The program control information signal generated by the operations center 202 provides data on the scheduling and description of programs. The program control information signal may be sent through the network controller 214 or, in an alternate configuration, directly to the set top terminal 220 for display to the subscriber. In the preferred embodiment, the program control information signal is stored and modified by the network controller 214 and sent to the set top terminal 220 in the form of a set top terminal control information stream (STTCIS). This configuration can accommodate, among other things, differences in individual cable systems and possible differences in set top terminal 220 devices.

10

15

20

25

The set top terminal 220 integrates either the program control signal or the set top terminal control information stream together with data stored in the memory of the set top terminal 220, to generate on-screen menu displays for assisting the subscriber in choosing programs for viewing. (Throughout the description the term "program control information" is being used to

indicate control information coming from the cable headend 208 to the set top terminal 220, whether it is sent directly from the operations center 202, processed by the network controller 214 and then forwarded to the set top box (STTCIS), or transmitted over telephone lines.)

The types of information that can be sent using the program control signal includes: number of program categories, names of program categories, what channels are assigned to a specific category (such as specialty channels), names of channels, names of programs on each channel, program start times, length of programs, description of programs, menu assignment for each program, pricing, whether there is a sample video clip for advertisement for the program, and any other program, menu or product information. In addition, the program control information signal may be used periodically to reprogram or reconfigure a set top terminal 220 or group of set top terminals 220 (described in detail in co-pending patent application Ser. No. 08/160,281, entitled, REPROGRAMMABLE TERMINAL FOR SUGGESTING PROGRAMS OFFERED ON A TELEVISION PROGRAM DELIVERY SYSTEM, filed by the same assignee incorporated herein by reference).

The goal of the menu driven program selection system 200 used with the present invention is to allow the subscriber to choose a program by touring through a series of menus utilizing a remote control 900 (Figure 3) or similar device providing cursor movement. The final choice

in the series of menus will identify one particular channel and one time for activation of that channel. Armed with a channel and activation time, the set top terminal 220 can display the selected program on the television for the viewer. To achieve this goal one embodiment of the present invention assigns an intelligent alpha-numeric code to each program. This alpha-numeric code identifies the category of the program, the menu in which the program should be displayed, its transmission time(s), and the position on the menu that the program should be displayed.

In this embodiment, the program control information, including menu codes, is sent continuously from the operations center 202 to the network controller 214, and ultimately to the set top terminal 220. For example, four hours worth of programming information can be sent via the program control information signal continuously using the information shown in Tables A-C.

Table A shows the basic programming information that may be sent to the set top terminal 220. The program descriptions shown are coded abbreviations. For example, C for comedy, N for news, S for sports, A for cartoons, and TX for text. If there is a textual description for a program, such as a movie, the description may be given following that program's coded description or may be communicated following the four hours' worth of programming information. As is shown in the coded listing, program descriptions for programs greater than a half hour in length

need not be repeated (each half hour). The video description code informs the set top terminal 220 of whether there is still or live video available to advertise the program.

5 For example, a sporting program may be assigned a code of B35-010194-1600-3.25-Michigan St. vs. USC. The letter B would assign the program to category B, sports. The second alpha-numeric character number 3 would assign the program to the third menu of the sports category. The third character of the code, number 5, assigns the program to the fifth program slot on the third menu. The next six characters, 01/01/94, represent the date. The following four characters, 1600 represent the start time which is followed by the length of the program and the program name. This entry represents a sports show, a college football game, which will be aired at 4:00PM on New Years day 1994.

10

15



TABLE A

12 PM					
	* Program name	* Program length	* Menu code	* Description	* Video
5	1 Cheers	.5	E24	C	N
	2 Terminator	2.0	A33	Tx	S
	3 PrimeTime	1.0	D14	N	N
	4 Football Special	.5	E24	S	N
10	*				
	*				
	*				
	*				
12:30 PM					
	* Program name	* Program length	* Menu code	* Description	* Video
15	1 Simpsons	.5	E14 & C13	C	S
	4 Football Game	5.0	B13	S	N
20	*				
	*				
	*				

25 In the 12:30 Channel 1 entry of Table A, two menu codes are shown. By allowing two menu codes, programs that may fit under two different category descriptions may be shown in both menus to the subscriber. With this minimal amount of information being communicated to the set top terminal 220 on a regular basis, the terminal is able to determine the proper menu location for each program and the proper time and channel to activate for the subscriber after his menu selection.

30 Table B shows an example Events Table that may be downloaded to a set top terminal 220 using the Event Data file, which contains information about events and

pricing. As shown in the table, the three columns of the Events Table identify the field number, the field itself and the type of information downloaded in the Event Data file. The first column contains the field numbers 1 through 11. The middle column contains the corresponding field parameters, including the event type, event ID, global channel ID, price, start time, end time, start date, end date, P- icon, name and description. The third column contains corresponding field type information. As shown in this field type information typically consists of an unsigned integer; hours, minutes and seconds; months, day and year; and ASCII character identifier.

TABLE B

Field #	Field	Type
1	Event Type	Unsigned Int
	1 = YCTV™ 2 = Pay-Per-View 3 = Reg. TV	
2	Event ID	Unsigned Int
3	Global Channel ID	Unsigned Int
4	Price (in Cents)	Unsigned Int

5	Start Time	H:MM:SS
6	End Time	H:MM:SS
7	Start Date	MM/DD/YY
8	End Date	MM/DD/YY
9	P-Icon	ASCIIZ
10	Name	ASCIIZ
11	Description	ASCIIZ

5

10

15

20

25

Table C shows an example Event Data file. In particular, Table C shows two data streams corresponding to two event types. The first data stream identifies a YCTV™ event in the first field. The second field designates the event ID, which is 1234 in this example. The third field includes the global channel ID number two. The fourth field indicates the cost of 50 cents for this event. The fifth and sixth fields indicate the respective start and end times of 3:00 a.m. to 3:00 p.m., respectively. The seventh and eighth fields show the corresponding start and end date, designated as 8/25/93 and 8/27/93, respectively. Field nine indicates the P icon set to PBS.PCX graphics file. Finally, fields ten and eleven indicate the name and description of the event selected, which in this case is Sesame Street and Barney. The second data stream in the Event.Dat example shown in Table C includes analogous information for

Terminator IV, which is designated in field one as a pay-per-view event.

TABLE C

Event Data Example

5	1 1234 2 50 03:00:00 15:00:00 08/25/93 08/27/93 phs.pcx Sesame Street & Barney's Sesame Street and Barney Abstract
	2 1234 2 50 20:00:00 22:00:00 08/25/93 08/25/93 14.pcx Terminator 4 Terminator 4 Abstract

10 The program control information signal and STTCIS can be formatted in a variety of ways and the on-screen menus can be produced using different methods. For instance, if the program control information signal carries no menu format information, the menu format for creating the menus can be fixed in ROM at the set top terminal 220. This method allows the program control information signal to carry less information but has the least flexibility since the menu formats cannot be

15

20 menu format information.

In the preferred embodiment, the menu format information is stored at the set top terminal 220 in temporary memory, either in a RAM or EPROM. This configuration provides the desired flexibility in the menu format while still limiting the amount of information needed to be communicated through the program control information signal. New menu format information would be sent using the program control information signal or

25

the STTCIS to the set top terminals 220 each time there was a change to a menu.

In the simplest embodiment, the menus remain fixed and only the text changes. Thus, the program control information signal can be limited to primarily text and a text generator can be employed in the set top terminal 220. This simple embodiment keeps the cost of the set top terminal 220 low and limits the bandwidth necessary for the program control information. Another simple embodiment uses a separate channel full-time (large bandwidth) just for the menu information.

4. Processing the Program Control Information Signal

Figures 6a and 6b show a more detailed schematic of the components of the cable headend 208, focusing on the interplay between the network controller 214 and the signal processor's 209 major hardware components. The network controller 214 uses, among other components, the signal processor 209 to implement its monitoring and control capabilities. Although the network controller 214 of the present invention will work with nearly any cable headend signal processing equipment, it is preferred that the signal processing equipment be modern equipment capable of handling digitally compressed video.

Figure 6a depicts an embodiment of the basic signal processing capabilities of the cable headend 208

and shows connections to components of the network controller 214. As shown in the figure, RF cable signals 205 are received at the headend 208 through a bank of integrated receiver demodulators (IRDs) 240. Each IRD 240 includes customary RF processing equipment, including a low noise amplifier, a demodulator and other filtering devices (not shown). As each RF feed is fed through the individual IRDs 240, the signals are manipulated and transferred to the demultiplexer and other signal processing equipment for further processing. The demultiplexer 242 splits each cable TV signal into its respective video and audio signal components. In addition, the demultiplexer 242 extracts data from the cable television signals and inputs such data to the control CPU 244.

The control CPU 244 exchanges control information with the network controller 214, as shown at 211. This control information is exchanged between the signal processor's control CPU 244 and the network controller CPU 224. In particular, the network controller 214 and signal processor 209 pass control information through the interface linking the two CPUs in order to perform any modifications to the program control information signal. The network controller CPU 224 oversees such modifications, accessing various network control databases 226 for guidance in instructing the signal processor's control CPU 244. The instructions provided

by the network controller 214 in turn guide the signal processor 209 in combining and/or adding programming signals and advertisements for transmission to the set top terminals 220.

5 The local insertion component 246 of the signal processor 209 allows the control CPU 244 to execute the instructions received from the network controller 214 and insert any local programming and advertisements. Once such regional programming and advertisements have
10 been inserted, the local insertion component 246 passes the various signals to a multiplexer 248 that combines the various programming and advertising signals. The output of the multiplexer 248 is transferred to RF modulator 250 that disseminates the composite video and audio signals to the set top terminals 220. The data
15 extracted from the cable television signals by the demultiplexer 242, which is also sent to the control CPU 244, is transmitted to the set top terminal 220 using a separate RF modulator 252.

20 The network controller 214 accommodates two-way RF data communications with the set top terminals 220. Upstream data transmissions from the set top terminals 220 are received by the network controller's control receiver 228. These upstream data transmission
25 capabilities are described in detail below.

Figure 6b diagrams another embodiment of a basic cable headend 208 having a network controller 214 and

more sophisticated signal processing equipment. Again, RF cable television signals 205 are fed into a bank of IRDs 240 as described above. These signals 205 are demultiplexed into individual video and audio signal components, with data being extracted and sent to the control CPU 244. The individual video and audio signal components are fed into a digital logic circuit 256 that is flexible enough to select individual video and audio signals for repackaging. The network controller 214 oversees such repackaging by: (i) receiving the program control information from the control CPU 244, (ii) modifying or manipulating the signal as necessary, and (iii) transferring the modified program control information signal back to the control CPU 244.

With instructions from the network controller 214, the control CPU 244 may insert local avails into the digital logic system 256 and execute the various selections of individual video and audio signals for subsequent transmission to the set top terminals 220. Once individual video and audio signals have been selected and all local insertions have been made, the outputs of the digital logic circuitry 256 are transferred to a serializer 258 which recombines all the signals into a serialized format. The serially-formatted signals are in turn transferred to RF modulators 250 for distribution over the cable network 200. The selection and recombining components of the signal processing

equipment are described in greater detail in a co-pending Patent Application, Ser. No. 08/160,283, entitled DIGITAL CABLE HEADEND FOR CABLE TELEVISION DELIVERY SYSTEM, incorporated herein by reference; however, such sophisticated combining circuitry is not necessary for the operation of the network controller 214. Rather, a simpler signal processing system may readily be used.

In the embodiments diagrammed in Figures 6a and 6b, the signal processor 209 may, acting alone or in conjunction with control instructions from the network controller 214, incorporate local programming and/or local advertisements into the program signals and forward the revised signal to the set top terminals 220. To accommodate this local programming availability, the signal processor 209 must combine the local signal in digital or analog form with the program signals 205 received from operations center 202. If a local cable system 200 uses a compression algorithm or standard that is different than the one used by the operations center 202, the signal processor 209 must also decompress and recompress incoming signals so they may be properly formatted for transmission to the set top terminals 220. In addition, the signal processor 209 performs any necessary signal decryption and/or encryption.

Figure 7 diagrams an alternative embodiment of a digital/analog cable headend 208. In particular, this

embodiment includes decompression and recompression capabilities, showing the types of signal processing components that the network controller 214 may control. As shown in Figure 7, the cable headend 208 receiver front-end, indicated at 260, demodulates the received transponder signals 205, which may contain four, six, eight or more audio/video channels of information, into a digital bit stream of multiplexed digitized MPEG or MPEG 2 format video. The signal processor 209 receives the multiplexed signals and initially performs any demultiplexing required to process the received signals. The demultiplexers 242 separate the multiplexed signals into separate individual MPEG or MPEG 2 format digital channels. Depending on the transponder signal received, the demultiplexer 242 may have four, six, eight or more cross connects to the combiner 264. The outputs of the demultiplexers 242 are selectively enabled by the control CPU 244. Those outputs of the multiplexer 248 that are enabled are then input to the combiner.

Decrypting may be necessary and can be conducted by a separate decrypting device 262 included as part of the signal processor's internal components. The signal processor's control CPU 244 may be controlled by a remote site (such as a national site) via a modem or similar connection 266. Therefore, the remote site is able to control the output of the demultiplexers 242. Alternatively, instead of enabling the outputs of the

demultiplexers 242, the inputs of the combiner 264 may be selected by the control CPU 244. By enabling or selecting multiplexer 248 outputs, the control CPU 244 is able to control which television programs are combined and transmitted to the viewers.

The combiner 264 combines the enabled or selected outputs of the demultiplexers 242 into the proper format and outputs the signals through a compressor 268, and an encryptor 270 (if desired), to a digital modulator 272. The modulator 272 outputs a modulated RF carrier combined with other carriers onto the cable distribution network 236. The set top converter terminals 220 in subscribers' homes select and demodulate a particular channel selected by the user. As selections are made, the set top terminal 220 stores the programs accessed in its local storage for later transmission to the network controller 214 at the cable headend 208.

5. Changing Menu Content by Modifying the Program Control Information Signal

Figures 8a through 8c are sample menu screens produced by a set top terminal 220 using the program control information signal. Figure 8a shows a menu which enables the viewer to select a program category from among a choice of eight program categories 1048. Figure 8b shows a menu 1050 for the viewer to select a hit movie from among ten hit movies 1052. Figure 8c

depicts a menu 1054 which provides information about a movie and enables a viewer to order the movie for viewing.

5 Figures 8a through 8c show text generated by a set top terminal 220. This text is generated using information received via the program control information signal by a text generator (not shown) in the set top terminal unit 220. Those portions of the text that generally remain unchanged for a period of weeks or months may be stored in EEPROM or other local storage. 10 For example, the text "HIT MOVIES from" 1056 will consistently appear on each hit movies' major menu. This text may be stored on EEPROM or other local storage. Further, text such as that which appears at the lower center part of the screen "PRESS HERE TO RETURN TO CABLE TV" 1058 appears many times throughout the menu sequence. This text may also be stored locally at the set top terminal 220. 15

20 Text which changes on a regular basis, such as the movie titles 1052 (or other program selections), will be transmitted to the set top terminal 220 by either the operations center 202 or the cable headend 208. In this manner, the cable headend 208 may change the program selections available on any menu by modifying the program control information signal sent by the 25 operations center 202 and transmitting the change.

It is preferred that the text, e.g., 1048, 1052, 1056, etc., be generated by the set top terminal 220 separately from the graphics because the text can be stored locally in a more compact manner requiring less storage space at the set top terminal 220. In addition, it allows for easy communication of text changes from the operations center 202 or cable headend 208 to the set top terminal 220.

Figures 8a through 8c show the use of day, date and time information 1060 on menus. This information may be obtained in a variety of ways. The day, date, and time information 1060 may be sent from the operations center 202, the cable headend 208 (signal processor 209 or network controller 214), the uplink site 204, or generated by the set top terminal unit 220 internally. Each manner of generating the day, date, and time information 1060 has advantages and disadvantages which may change given the particular embodiment and costs.

In the preferred embodiment, the day, date, and time 1060 are generated at a central location such as the operations center 202 and are adjusted for regional changes in time at the cable headend 208. In particular, the network controller 214 modifies the PCI signal to accommodate regional day, date and time information and changes and additions in regional programming and advertisements. These modifications are automatically



processed by the network controller CPU 224 upon initiation of the Modifying PCI software routine, as described below. In an alternate embodiment, the network controller's control station operator can manually enter programming, advertising and menu modifications.

6. Receiving Information from Set top Terminals

10 The network controller 214 is equipped to receive information from the set top terminals 220 on a regular or random basis. Figures 9a and 9b diagram separate embodiments for upstream data transmission for a digital/analog cable headend 208. In particular, Figure 15 9a diagrams an out-of-band two-way data transmission system 280 wherein satellite feeds 262 are received at the cable headend 208 by a number of satellite receivers 284 and digital signal processing equipment 286. The 20 satellite receivers 284 are used for analog transmissions and the digital signal processing equipment 286 is used to process digital programming signals. The analog signal paths allow analog cable television programming signals to be received by the set 25 of satellite receivers 284 and to be passed to a series of modulators and scramblers 288 the output of the modulators and scramblers 288 is sent to an RF combiner 290.

A data transmitter (Data Tx) 292 makes use of the control information transferred to the signal processing equipment from the network controller 214. This data transmitter inserts data into the RF combiner 290.

5 Through the use of a separate data transmitter, any downstream data transmissions may be sent to a set top terminal 220 on an out-of-band frequency (i.e., out of the frequency band used for video signal transmissions).

10 Digital signals are also input to the RF combiner 290 from the digital signal processing equipment 286. These digital signals are typically assigned to separate frequency bands. Once the data, analog and digital signals have been combined using the RF combiner 290, the composite signals are further processed at the cable headend 208 for distribution over the cable network. This further processing involves using a diplex filter 294 that accommodates two-way RF communications over the cable distribution network.

15 The diplex filter 294 requires that the various sets of signals be translated to different frequency bands. Typically, services to the home are sent in a downstream band, which begins at 54 MHz and extends today to typically 550 MHz. Other systems that use a maximum frequency less than or greater than 550 MHz, however, 20 may readily be accommodated by the embodiment shown in Figure 9a. Downstream services may include TV 25

channels, FM radio, digital/audio signals and various control and information data streams.

Upstream transmissions from the set top terminal 220 are typically sent in the frequency band between 5 and 50 MHz. Other frequency limits may, however, be employed in special cases. For example, the industry is currently experiencing movement toward using 5 to 42 MHz for upstream services.

Although the diplex filter 294 is not an inherently bi-directional device, it may be made bi-directional by splitting the spectrum between downstream and upstream signals, as described above. The diplex filter 294 effectively becomes bi-directional by passing high-band signals in the downstream direction and passing low-band signals in the upstream direction. For downstream transmission capability, all signals in the high-band of 50 to 550 MHz are passed to a fiber/coax translation point, indicated generally at 300.

At the fiber/coax translation point 300, optical energy is relayed to the various optical nodes 304. This distribution of optical energy typically involves splitting the optical energy among the nodes 304 and transporting the energy downstream on one or more downstream fibers. In addition, electrical energy signals are sent over coaxial cables, through a series of amplifiers 308 along the cable for distribution to individual subscribers. Individual subscribers simply tap into the amplifiers

along the coaxial cable in order to receive programming and downstream data signals.

Upstream data transmission are sent to the cable headend 208 from each optical node 300 over fiber and input into the cable headend's RF combiner 308. Upstream transmissions over cable are accommodated using carrier frequencies in the lower frequency band. These upstream data transmissions over the coaxial cable are passed through the diplex filter 294, which filters out all high-band frequencies and passes all low-band frequencies. Subsequently, the diplex filter 294 transfers such low-band frequencies to the RF combiner 308. The RF combiner 308 combines all upstream data transmissions from the set top terminals 220 and inputs these combined data signals into the network controller 214 for later processing.

Figure 9b shows an alternative embodiment to Figure 9a. In particular, Figure 9b shows the same overall configuration as the embodiment above (and is commonly numbered) although downstream data transmissions from headend 208 to the set top terminals 220 are accomplished through in-band two-way data transmission. Thus, the primary difference between the diagrams shown in Figures 9a and 9b is that the latter embodiment uses a method of inserting data into the downstream programming signals themselves for

distribution to the set top terminals 220 in the cable network.

5 Basically, the data placed on the programming signals using a set of data inserters 312 that are electrically connected to each modulator and scrambler component 258. In this way, data can be inserted in-band along with video and audio signals, thereby modulating the data on the same respective carrier frequencies used by the video and audio signals. The inserted data is thus
10 combined with video and audio signals and input into the RF combiner 290 for downstream distribution. As described above, digital signals are also combined using the RF combiner 290 and disseminated over the cable network. Upstream transmissions are accomplished as
15 described above in conjunction with the discussion for Figure 9a.

Upstream information received from the set top terminals 220 typically includes, for example, program access data gathered at each set top terminal 220. Such
20 information may be communicated to the network controller 214 through a variety of methods including any of the following methods: (1) cyclic polling, (2) random access, and (3) telephone modems. Cyclic polling and random access methods make use of the two-way RF system diagrammed in Figures 9a and 9b, described
25 above.

As described below, the preferred embodiment employs a cyclic polling method. Although various polling schemes will work with the present invention, a roll-call polling scheme is preferred over other schemes such as hub polling or token-passing since roll-call polling provides the greatest degree of centralized control.

Using this preferred method, program access information is stored at each set top terminal 220 until it is polled by the network controller 214 for information retrieval using a polling request message format 920 as shown in Figure 10a. This frame format 920 may include such program control information as shown in Tables A-C above, typically consisting of six fields: (1) a leading flag 922 at the beginning of the message, (2) an address field 924, (3) a subscriber region designation 926, (4) a set top terminal identifier 928 that includes a polling command/response (or P/F) bit 930, (5) an information field 932, and (6) a trailing flag 934 at the end of the message.

The eight-bit flag sequence that appears at the beginning and end of a frame, 922 and 934, respectively, is used to establish and maintain synchronization. Such a sequence typically consists of a "01111110" bit-stream. The address field 924 designates a 4-bit address for a given set top terminal 220. The subscriber region designation 926 is a 4-bit field that indicates the

geographical region in which the subscriber's set top terminal 220 is housed. The set top terminal identifier 928 is a 16-bit field that uniquely identifies each set top terminal 220 with a 15-bit designation followed by an appended P/F bit 930. Although field size is provided by this example, a variety of sizes can be used with the present invention.

The P/F bit 930 is used to command a polling response from the set top terminal 220 addressed, as described below. The frame format 920 also provides a variable-length information field 932 for other data transmissions, such as information on system updates. The frame format 920 ends with an 8-bit flag 934 (or trailing flag) that is identical in format to the leading flag 922, as set forth above. Other frame formats will be apparent to one skilled in the art and can be easily adapted for use with the system.

Using any such polling request message format 920, the network controller 214 interrogates each set top terminal 220 sequentially, one by one. In this type of access strategy, the network controller 214 is designated as the central controller of the cable distribution network 200 and is responsible for control of the communications links between itself and the set top terminals 220. This control includes issuing commands to the set top terminals 220 and receiving responses back from the set top terminals 220.

Basically, the network controller 214 instructs the signal processor 209 to transmit to each set top terminal 220 a polling request, which asks whether a set top terminal 220 has any information to transmit. The set top terminals 220 are identified by the unique address and set top terminal identifier 928. It is preferred that the set top terminal 220 transmit information and messages to the network controller 214 only when given permission by the network controller 214 to do so.

Where, for example, specialty programs have been accessed since the previous poll, the set top terminal 220 is given permission to transmit a polling response in the form of a status report that includes any such access information. The network controller's control receiver 228 is tasked with the receipt of set top terminal 220 polling responses or status reports. These status reports generally include information that allows the network controller 214 to track a subscriber's program access history. As described above, the control receiver can store the status reports locally and/or transfer them to the network controller CPU 224.

The network controller CPU 224 immediately processes each polling response as it is received from each set top terminal 220. The network controller CPU 224 updates pertinent databases 226 with the received information, and then sends another polling request to

the next set top terminal 220 on its list. A set top terminal 220 with no information to transmit so indicates in a reply to the network controller 214. Once all set top terminals 220 have been given permission to transmit status reports, a cycle is complete and a new cycle begins.

Through a polling cycle, the network controller 214 acquires the information needed to operate the system 200. During the cycle, the network controller 214 sends signals to the set top terminals 220 to authorize both their operation and access to specific channels. If, for example, a subscriber has failed to pay a recent bill, the network controller 214 can deauthorize the subscriber's set top terminal 220. Likewise, when a subscriber orders a program or channel, the network controller 214 checks the subscriber's account for good standing by reading the proper database file. After the check, the network controller 214 then either authorizes or deauthorizes access by the set top terminal 220 using the data transmitted in a modified program control information signal. As a result, the cycle requires a series of requests and responses to operate.

Figure 10b shows an example frame format 920' for the status reports received from the set top terminals 220 during the polling cycle. This frame format is substantially identical to the polling request message format 920 (Figure 10a), and includes: (1) a

leading flag at the beginning of the message, (2) an address field, (3) a subscriber region designation, (4) a set top terminal identifier that includes a polling command/response (or P/F) bit, (5) an information field, and (6) a trailing flag at the end of the message, each designated by a common number with respect to Figure 10a, but with the prime indicator (') added.

Again, the information field 932' remains variable in length so that the status of an indeterminate number of programs accessed, as represented at 933', can be included in the frame. In this way, the control message length of the polling request message is minimal since the network controller 214 does not transmit such access information. After a polling response by a given set top terminal 220, however, the control message length increases in proportion to the number of programs accessed.

During transmission, the P/F bit 930, 930' is used to carry out the polling function. In particular, the P/F bit 930 is set to a "1" position to command a polling response from the set top terminal 220 whose address is identified in the frame 928. The set top terminal 220 addressed must respond to the command with the same P/F bit 930' also set to the "1" position. The response will include the number of programs accessed and their corresponding event identification numbers as shown in Figure 10b at 933'. In cases where the set top terminal

220 has not accessed any programs since the previous polling cycle, the set top terminal 220 responds with the P/F bit 930' set to "1" and the programs access block denoting zero programs accessed.

5 The second method for the network controller 214 to receive information from the set top terminals 220 is through the use of a random access scheme. In an alternate embodiment that uses this method, individual set top terminals 220 can send control-related messages
10 to the network controller 214 without being polled. This scheme is particularly useful in networks where subscriber regions include potentially large numbers of subscribers. High concentrations of subscribers may be found, for example, in large metropolitan areas. In such
15 cases, the polling cycle can be replaced with a more sophisticated random access strategy such as carrier-sense multiple access with collision detection (CSMA/CD). In this scheme, each set top terminal 220
20 must "listen" before it transmits and then does so only if it senses an idle medium. When the return link to the network controller 214 is silent, a given set top terminal 220 can transmit its messages. Any messages sent from a set top terminal 220 to the network controller 214 would set the P/F bit 930' to a "0" position to indicate
25 that the message is not in response to any command or polling request. In addition to CSMA/CD, other random

access schemes can be used with the system, such as CDSL.

The third method for the network controller 214 to receive information from the set top terminals 220 is through the use of telephone modems. In an alternate embodiment, the set top terminals 220 communicate program access information and orders to the network controller 214 using telephone modems. In this embodiment, the set top terminals 220 are equipped with a modem port to facilitate such operation. Thus, communications between a given set top terminal 220 and the network controller 214 can be established over telephone lines when cable traffic or other primary traffic is congested. The preferred method of using telephone modems is in combination with a control or "hit" signal from the network controller 214. A group (or region) of set top terminals 220 is "hit" simultaneously by the network controller 214 via the cable. Only those set top terminals 220 within the group that have data for the network controller 214 call the network controller 214 by modem. The network controller 214 is equipped with a bank of modems (organized to roll-over telephone calls) to answer the incoming calls.

Among the three methods discussed for the network controller 214 to receive information from the set top terminals 220, the use of the cyclic polling

scheme depicted in Figures 10a and 10b, is preferred. Polling is preferred because it allows the network controller 214 to conduct and control communications with set top terminals 220 over the cable network in an orderly fashion. In particular, the network controller 214 can schedule data retrieval by polling the set top terminals 220 one by one. A random access method, on the other hand, does not allow the network controller 214 to maintain such orderly communications. Instead, the network controller 214 receives data from the set top terminals 220 at random, depending on when the cable medium is idle. This random reception of data lessens the degree of control that the network controller 214 has over set top terminal transmissions. Likewise, the third method, which uses telephone modems, is less desirable than the polling method since the use of modems does not allow for upstream interactivity over the cable medium.

7. Processing Information Received from Set top Terminals

Regardless of the scheme used by the set top terminals 220 to access the network controller 214, any polling responses and upstream interactivity is received by the network controller's control receiver 228 as shown in Figure 11, depicting the components of the control receiver 228, which includes a demodulator 310 and demultiplexer 313 to demodulate and demultiplex

transmissions received from any set top terminal 220 in the cable distribution network 200. As described above, the control receiver 228 transfers, through a control buffer 315, the received information to the network controller CPU 224 for processing.

Processing is accomplished by the network controller CPU 224. Operator instructions are input to the network controller CPU 224 through the operator control station 234 that includes, for example, a computer/workstation with a CRT display, printer and other peripherals. Multiple operator control stations 234 can be used to assist in control operations.

Regional operator control stations (not specifically shown, but substantially identical to stations 234) may be used and may include multiple operator control stations each assigned to a particular subscriber region corresponding to a geographic region where set top terminals 220 are located. Thus, each regional operator control station is assigned to a subscriber region, providing monitoring and control capabilities over such regions. All regional program control information is transferred to the network controller CPU 224 for processing, as in the case where a single control station 234 is used. Likewise, during this processing, portions of the network control databases 226 may also be updated.

No set number of databases 226 are required for the network controller 214 to perform its operations, and a single temporary database may be used. In the preferred embodiment, however, the network controller 214 uses several databases (indicated at 226) that are accessed during network control operations. These databases 226 are identified in Figure 11 and include: (1) the Viewer Profile database 314, (2) the Account/Billing database 316, (3) the Program Library database 318, (4) the Program Scheduling database 320, (5) the Advertisement Library database 322, and (6) the Advertisement Scheduling database 324.

Figure 12 shows one example of a network controller's basic database structure including the databases identified in the preceding paragraph. The data stored in these databases is not simply raw data. Rather data may be processed, correlated and appropriately indexed to create a true relational database 226.

As shown in Figure 12, the Viewer Profile database 314 includes: (i) a Set top ID File, (ii) a Subscriber Region File, (iii) a Customer ID File and (iv) a Viewer Log File, the latter three files being indicated generally as a file group 332. The Set top ID File 330, common to each of the databases comprising the network controller's database 226, contains set top converter records with each record representing a unique set top terminal 220. Examples of information stored in this file includes set

top terminal type, software version and set top terminal identification/serial number. The Set top ID File 330 contains the key data that links each relational database with one another, as described below.

5 The Subscriber Region File, part of file group 332, includes information such as headend 208 assignment, regional operator control workstation assignment and a designation for the subscriber's geographical area. The Customer ID and Viewer Log Files, part of file group 332, 10 include the subscriber's personal information, such as name, address and telephone number, and information on the subscriptions to cable services for each customer as well as a personal profile for each viewer, respectively.

 The personal profile consists of demographic 15 information that may be gathered in a number of ways. The set top terminal 220 builds the personal profile for each viewer and stores the information in a memory file by viewer name. To build a personal profile in the preferred system, the viewer answers a series of 20 questions presented on a series of menu screens. These personal profile screens request the viewer to input information such as name, sex, age, place of birth, place of lower school education, employment type, level of education, amount of television program viewing per 25 week, and the number of shows in particular categories that the viewer watches in a given week such as, sports, movies, documentaries, sitcoms, etc. Any demographic

information which will assist the set top terminal 220 in targeting advertisements to the viewer may be used.

In addition to gathering demographics at the set top terminal 220, the personal profile can be compiled using other methods. For instance, the information can be gathered using questionnaires sent by mail and subsequently entered in the Viewer Profile Database 314 by the network controller's control station operator.

As an alternative to gathering demographic data, a simulated profile can be generated using an algorithm similar to that described below that analyzes access history and viewing habits. Using test information generated from a statistically significant number of viewers, the simulated profile algorithm estimates the viewer's age, education, sex and other relevant information. The analysis requires reviewing the viewer's programs watched and statistically comparing the viewer's programs watched with the test group. Also, the algorithm can place the subscriber or viewer in a viewer category. This analysis is transparent from the subscriber's point of view and attempts to accurately profile the viewer. Various viewers or viewer categories can later be targeted with different advertisements.

The Account/Billing database 316 includes (i) the Set top ID File 330, and (ii) an Account History File, and (iii) a Billing File, the latter two files indicated at 338. The Set top ID File, as described above, contains

information unique to each subscriber, including set top terminal type, software version and set top terminal identification/serial number. The Account History and Billing Files contain information concerning each subscriber's past bills and account record and information on the most recent bill, including data from which the next billing report can be generated, respectively.

The Program Library database 318 include (i) the Set top ID File 330, and (ii) a Programs File, (iii) a Preview File, (iv) a Program Category File, (v) a Price Category File and (vi) Service File, the latter five files identified at 344. As usual, the Set top ID File identifies each set top terminal 220 by identification number. The Programs File contains information on every program offering in the system, including name, length and type of program. The Preview File contains information on previews for specialty programs stored in the Programs File. The Program Category File contains a set of categories into which each program may be placed, such as movies, sports, science fiction and news. The Price Category File contains information on pricing for various categories of programs, grouping programs and services into categories by price. The Service File maintains information on the various cable services available in the system 200.

The Program Scheduling database 320 includes (i) the Set top ID File 330, and (ii) an Access History File, (iii) a Programs Watched Matrices File and (iv) a Program Scheduling Library, the latter three files indicated at 350. The Access History File contains information on the programs that the set top terminal 220 has accessed and the Programs Watched Matrices contains information on the number of programs watched in a given program category during different times of day. Relative to the Programs Watched Matrices file, a programs watched matrix is shown in Fig. 16 and further described below. The Program Scheduling File contains information on the times of day and the corresponding programs that are being offered for viewing at each subscriber location.

The Advertisement Library database 322 includes (i) the Set top ID File 330, and (ii) an Advertisements File, and (iii) an Advertisement Category File, the latter two files being indicated at 354. The Advertisements File contains information on every advertisement in the system, including name, length and type of advertisement, and the Advertisement Category File contains a set of categories into which each advertisement can be placed.

The Advertisement Scheduling database 324 includes (i) the Set top ID File 330, and (ii) an Advertisement Selection File, and (iii) an Advertisement Targeting File, the latter two files identified at 358. The Advertisement Selection File contains information

on the advertisements that have been offered to each subscriber and keeps track of the ones that have been selected. The Advertisement Targeting File contains information on the advertisements and advertisement categories that have been chosen by the system as being
5 of the most interest to a specific subscriber.

The network control databases 314, 316, 318, 320, 322, 324 comprising the database 226 are relational databases generally keyed to information in a single file. Specifically, the relational key is a set top terminal 220
10 identification number stored in Set top Terminal ID File 330, as shown in Figure 11. This set top terminal identification number allows the database files that correspond to a particular subscriber to be linked together by a common reference. In other words, the databases are structured such that subscribers are
15 referenced in each database file by a unique set top terminal identification number. In this way, each database may be accessed based on set top terminal identification number alone. Thus, using a subscriber's set top terminal identification number, the network controller CPU 224 can access and process information pertaining to that subscriber from any of the above
20 described database files. In configurations where multiple set top terminals 220 are allocated to a single customer (or household), a unique subscriber identification number may be added to the database 226
25

to group the set top terminals 220 by customer. With the set top terminal identification as a relational key, many additional databases may be created that correlate and store pieces of subscriber-specific information from the six databases and underlying files.

6. Overview of Software Routines

Figure 13 shows the major software routines initiated and executed by the network controller CPU 224. These routines are: (1) the Modifying PCI routine 370, (2) the Polling Cycle routine 372, (3) the Advertisement Targeting routine, and (4) the Account/Billing routine 376. Together, these routines, along with the operator entry and update functions 380, 382, respectively, enable the network controller 214 to perform its major functions.

The Modifying PCI routine 370 is the software that enables the network controller 214 to modify the program control information (PCI) signal received from the signal processor 209. This software routine generally allows the network controller CPU 224 to modify the PCI signal content so that changes and additions in programming and advertisements can be accommodated. Such changes and additions include access authorizations and deauthorizations in the form of authorization and deauthorization messages, respectively.

The Polling Cycle routine 372 is the software sequence that interactively executes the network controller's polling cycle allowing the network controller 214 to schedule and perform polling of all set top terminals 220 operating in the system 200. The software also provides the network controller 214 with a means of processing status reports received from set top terminals 220 in response to polling requests. For a random access system (not depicted), the software of this routine 372 would be changed.

The Advertisement Targeting routine 374 is the software that generates packages of television commercials and advertisements geared towards particular viewers and makes use of a viewer's demographic information and viewing habits to determine those advertisements that are of most interest to that particular viewer. In so doing, the routine 374 outputs packages of advertisements targeted towards each viewer.

The Account/Billing routine 376 is the software that the network controller CPU 224 runs to generate billing reports for each set top terminal 220. In general, the routine 376 correlates the programs accessed with pricing information to generate each report.

9. Modifying PCI Routine

Figure 14 shows a software flow diagram for the network controller's Modifying PCI routine 370. The

Modifying PCI routine (or sequence) is initiated, block 384, automatically by the network controller CPU 224 upon receipt of the program control information (PCI) signal from the signal processor 209. Once the network controller 214 receives the PCI signal, the network controller CPU 224 begins processing the signal by reading the PCI data carried by the signal, block 386.

After reading the PCI data, the network controller CPU 224 "calls" other routines to interactively process data and continue the modification process for each set top terminal 220. First, the network controller CPU 224 calls the Polling Cycle routine 372, at block 388, in order to request data retrieval of the information stored at individual set top terminals 220. Such information includes data on the programs accessed and those ordered for later viewing. As polling responses are received from the set top terminals 220, the network controller CPU 224 next calls, block 390, the Advertisement Targeting routine 374, which generally arranges groupings of commercials for different subscribers based, in part, on viewer demographic information and program access history.

The network controller CPU 224 next calls 392 the Account/Billing routine to begin processing all programming and channel access requests. The Account/Billing routine determines, among other things, whether the subscriber's account is in good standing.

verifying that past bills have been paid and that access authorization is warranted. Upon completion of this verification process, a verification message will be sent to the network controller's operator control station 234
5 indicating that access should be granted.

In the preferred embodiment, an access authorization code may automatically be processed by the network controller CPU 224 and appended to the PCI signal originally received from the signal processor 209.
10 This modified PCI signal and access authorization code will then be transferred back to the signal processor 209 for transmission to the set top terminals 220.

With continued reference to Figure 14, in an alternate embodiment that uses the Modifying PCI Routine 370, at blocks 394 and 396, the operator
15 manually enters any changes in programming and menu content, along with access authorizations, into the program scheduling database 320. The manual entry of programming and menu content in this embodiment,
20 blocks 394, 396, requires that the operator access the database information generated and updated by the other routines and make necessary changes in the program scheduling database. The network controller CPU 224 reads this updated database information, generates a
25 modified PCI signal, and sends, block 398, the signal to the signal processor 209.

If a subscriber account is delinquent, access to any new programs or channels ordered will not be authorized. Instead, the network controller CPU 224 will deny authorization and generate a deauthorization message to be included in the PCI signal that will be returned to the signal processor 209 for transmission to the set top terminals 220. Alternatively, the network controller CPU 224 generates a delinquency message that is transferred to the CRT display at the network controller's operator control station 234. Upon reviewing the message, the operator may then manually enter message text to be included in the PCI signal that informs the subscriber of a delinquent account.

10. Polling Cycle Routine

Figure 15 shows a software flow diagram for the network controller's Polling Cycle routine 372, which iteratively executes the network controller's polling cycle. The number of iterations correspond to the number of set top terminals 220 being polled. The network controller CPU 224 initiates the Polling Cycle sequence periodically on a predetermined basis, block 400. Typically, this period is set by the operator at the network controller's operator control station 234 at once per day, although other periods (e.g., multiple times per day or once per week) can be used.


Upon initiation of the sequence 400, as depicted at function block 402, the network controller CPU 224 reads

the Set top Terminal ID File 330 and begins generating, block 404, a polling request frame (shown in Figure 10a and described herein above) for the first set top terminal 220 identified in the file 330. Once the necessary polling request information is complete, the frame is transferred to the signal processor CPU 244 through the interface between the signal processor 209 and network controller 214. After transfer to the signal processor 209, the frames may be transmitted to the set top terminals 220, block 406. Meanwhile, the network controller's control receiver 228 awaits the corresponding response.

Upon receipt of a polling response, as depicted at block 408, the network controller CPU 224 reads the received information from the control buffer 315. The network controller 214 reads the information field of the polling response frame format, as described above. The network controller CPU 224 processes, indexes and stores the data in an appropriate format, updating the corresponding database files with the information received, block 410. The processing and indexing of the raw data into a relational database 226 is important to the ability of the network controller 214 to quickly take actions such as targeting commercials without lengthy processing time. The polling routine subsequently returns to the Set Top Terminal ID File 330, as shown at decision block 412, to continue the polling cycle for the

next set top terminal 220 identified in the file 330. When the routine 372 sequences through the last set top terminal 220, the cycle is complete and the routine 372 ceases until the next polling period.

5 Most often, the files that require updates during the polling cycle are the Access History File and the Programs Watched Matrices File, both indicated generally at 350 in Figure 12, and the Account History File 338. For example, Figure 15 shows an example of a 30-day
10 programs watched matrix, denoted 351, for one set top terminal 220 (not shown in Figure 16). The matrix 351 is divided into six rows, corresponding to six four-hour time slots. The columns of the matrix 351 are divided, as necessary, by the program categories available for
15 viewing. Each entry in the matrix 351 denotes the number of programs watched in a particular program category and time period.



After the status report is received on each set top terminal 220, the polling response routine (see Figures
20 10a and 10b) determines which time slot and category of program numbers in the matrix 351 need to be increased. Thus, entries in the matrix 351 are updated upon receipt of each set top terminal's polling status report, thereby maintaining a running total of the programs watched. For
25 example, during the 0600-1200 time period, the matrix 351 shows that this set top terminal 220 has been used to watch ten movies during the past month. Preferably

the program watched identifying information is stored in addition to the running totals in the Programs Watched Matrices file. Use of programs watched matrices is further described in the following section describing the Advertisement Targeting routine.

11. Basic Advertisement Targeting Routine

Figure 17 shows the seven primary functions of the basic advertisement targeting routine 374. The function of this routine is to target video for set top terminals 220 based on historical viewing data and other data that is available at the network controller 214. Advertisements that may be targeted include video, commercials and infomercials, with infomercials being time varying video segments (e.g., thirty seconds, fifteen minutes).

When initiated, block 420, the first subroutine, identified at function block 422, accesses the programs watched matrices (exemplified by matrix 351) stored in the Programs Watched Matrices file in the Program Scheduling database 320. The subroutine uses a unique set top terminal ID to access a specific matrix for one set top terminal 220. These matrices are maintained and updated by the polling response routine.

The second subroutine, function block 424, which develops other matrices based on other available information, is an optional subroutine not required for

the functioning of the system. For groups of set top terminals 220 or for each individual set top terminal 220, matrices may be developed based on the demographic information, billing information, pricing information, age information and other information which may be stored in the network controller 214 databases.

The third subroutine, block 428, processes all matrices through a set of correlation algorithms. In particular, this subroutine 426 takes matrices developed in the first two subroutines and processes the matrices until reaching a final matrix.

Figure 18 diagrams an embodiment of this matrices processing subroutine 426 which is called by the advertisement targeting sequence shown in Figure 17. As shown in Figure 18, the subroutine 426 is initiated 427 and then accesses or queries, block 428, the programs watched file and gathers information regarding either an individual subscriber or a node of subscribers. The software can gather the programs watched information in this way for individual subscribers or a set of subscribers.

Once the programs watched information has been gathered from the databases, the routine 426 selects and groups, function block 430, programs watched based on program categories and time slots. The software initially takes each program category (e.g., sports, news,

movies, etc.) and establishes the number of programs watched for a given time slot. The time slots may be set to any length of time, including, for example, one, two, three or four hour timeframes. The software will loop through such a counting process for each group and timeslot and then proceed to build a programs watched matrix, block 432, based on the program categories and time slots. Essentially, all programs watched in a particular category and time slot will be entered into the programs watched matrix. Once the matrix has been built, the subroutine 426 will process the matrix for a given subscriber or node of subscribers through the correlation algorithms.

A number of correlation algorithms may be used to weight each selected program category group. For example, as shown at block 434, a sum of squares algorithm may be used to determine the weighting. Once the groups have been weighted, the weighted groups will be correlated, as at block 436, with various advertisements stored in the network control databases. The software can then select a set of the most heavily weighted advertisements for transmission to individual subscribers or sets of subscribers in a cable distribution network node. Having determined the weightings of each group and prioritizing the groups accordingly, the subroutine returns 438 to the advertisement targeting sequence 374 of Figure 17.

Referring back to Figure 17, the fourth subroutine, as represented at function block 428, uses the final matrix developed by the correlation and weighing algorithm described above, to select a grouping (or selective filter) for each set top terminal 220. The final groupings of advertisement that may be sent to the set top terminals 220 or node of set top terminals 220 may use a subroutine as diagramed in Figure 19.

The subroutine 428 depicted in Figure 19 is called or initiated by the advertisement targeting sequence 374 of Figure 17 in order to determine the final groupings. Basically, this subroutine selects a set of commercials that will be used in the chosen groupings, function block 444. This selection process typically involves advertisements from various advertisement categories (from a number of advertisers which have purchased "air time"). Each advertisement will subsequently be assigned a number of times that it will be shown in a given timeframe, block 446. This frequency of display may be based on various factors, including the number of requests and cost paid by the respective advertisers to have the commercial displayed. Such factors are used in the next step of the subroutine, block 448, which assigns a weighting to specific commercials or advertisements in each advertisement category or group. These weightings are used to prioritize the advertisements

that will be sent to individual set top terminals 220 or nodes of set top terminals 220.

Once the advertisements have been weighted, the software executes its correlation algorithm, 450, using selected criteria (i.e., the various factors used to weight the advertisements) as well as the output of each programs watched matrix. Any number of correlation algorithms and weighting algorithms may be used with the software, including the sum of squares weighting algorithm described above.

The results from the correlation algorithm subsequently determine the advertisements and programming material that is sent to the signal processor 209 for distribution over the cable network, as represented at block 452. Once the subroutine 428 completes these steps, the network controller CPU 224 updates the account and billing database based on the ads that are sent to the signal processor 209 for subscriber viewing, as shown at block 454. These billing database updates allow the advertisers to track the costs and frequency of the advertisements targeted to specific set top terminals 220 or nodes of set top terminals 220. Following the updates, the subroutine returns to the advertisement targeting sequence shown in Figure 17, block 455.

Referring to Figure 20a, set top groupings (A through E) 460 are shown. The number of set top

groupings available is determined by the bandwidth available to transmit commercials. The bandwidth of the system will limit the number of commercials which are available at the set top terminal 220 at any given time.

5 Referring back to Figure 17, the fifth subroutine, represented at function block 466, prepares set top group information for transmission to the set top terminals 220. This subroutine 466 modifies the PCI signal and
 10 includes set top group information in the information field of the frame format given earlier. The various methods for transmitting the group information to the set top terminals 220 are described below.

The sixth subroutine, block 468, selects the target video and is the last decision making process in targeting
 15 a commercial for a viewer and, can be performed by either the set top terminal 220 or the network controller 214. In the preferred embodiment, the set top terminal 220 performs this last step by correlating (or matching)
 20 the program being watched by the viewer with the set top group information that has been previously transmitted by the network controller 214, and the targeted video is then displayed, as shown at block 470. Figure 20a shows an exemplary table matching set top
 25 terminal groups 460 and program category being watched 470 with a specific channel (continuously) showing commercials. The commercial channels are shown in Figure 20b at 474 and are assigned Roman numerals I

through X, for example. The number of set top groupings and channels showing commercials can vary. Figure 20b shows a division of available bandwidth to carry ten videos, ten commercial channels. In this example, the channels 474 are numbered 101-110.

The network controller 214 will transmit group information to a set top terminal shown as row names 460 on Figure 20a. The network controller 214 will also transmit data which informs the set top terminal 220 which of the multiple commercial channels 474 is assigned to a television program category shown as Column 470 on Figure 20a. Each set top terminal 220 only requires the data related to that set top terminal's assigned group (or row). For example, in Figure 20a, the set top terminal in group A (row A) is provided with data on the commercial channel which are assigned for sports programs as I, children's programs as IV and movie category as III. In this manner, each set top terminal 220 is only required to store information related to its own grouping. Therefore, a set top terminal 220 which is in group A only needs to store the information related to group A, which is found in row A of Fig. 20a. This information includes one commercial channel assignment for each of the eight program categories. Using this information, the set top terminal 220 first determines the category of the television program currently being watched and then is able to quickly determine which



5
10
15
20
25

channel to switch the viewer when an advertisement availability occurs during the program.

The network controller 214 can also perform the step of correlating program category watched 470 and set top terminal grouping 460 to select the target video. In order for the network controller 214 to perform this function, it must have information on the program currently being watched by the viewer. To obtain this information in a polling system, set top polling must occur on a real-time basis (i.e., 10 minutes).

During the target commercial selection process, the set top terminal programming will default to the existing commercial during a program if it is missing any of the information needed to determine which of the continuously playing commercial channels to show. In alternative embodiments, the default that is shown on the regular programming channel will correlate with one of the assigned set top groupings and program categories. Figure 20a shows, at 478, that the default has been assigned to set top terminal grouping C for program categories "children" and "entertainment."

The three preferred methods to transmit targeted commercials to a set top terminal 220 are: (1) the Additional Bandwidth method (or individual video access); (2) the Multiple Channel method, and (3) the Split Screen method. Each method has certain advantages and disadvantages. The Additional Bandwidth method

allows the most flexibility by more specifically
targeting commercials before the commercials are
transmitted to a set top terminal 220. However, it
requires a great deal of available bandwidth in the
5 delivery system. This is difficult with a cable system
200 but possible when a telephone or personal
communications system is used to transmit the
commercials to the set top terminal 220.

The Additional Bandwidth method allows the
10 network controller 214 to run through a set top
terminal's specific correlation algorithms and target
specific commercials from hundreds for each set top
terminal 220. This method allows for the greatest
customizing of targeting and allows for a greater
15 selection of commercials to be shown. Only after a
commercial advertisement is selected by the network
controller 214 for the specific set top terminal 220 does
transmission of the commercial occur.

The Multiple Channel method requires a set top
20 terminal 220 "transparently" to change channels during a
scheduled advertisement from the channel of the
currently viewed program to the channel which is
carrying the targeted commercial. Although this channel
changing method may be transparent to the viewer, it
25 creates difficulty in terms of timing and synchronizing
the commercials to begin and end during an
advertisement availability occurring in the normally

scheduled program. The channel changing is done within the set top terminal 220 using the existing tuner(s) (not depicted). Alternatively, in set top terminals 220 equipped with two tuners, the terminal can use the second tuner to tune the channel showing the commercial. (Set top terminals with two tuners are described in detail in co-pending patent application, Ser. No.08/160,194, entitled, ADVANCED SET TOP TERMINAL FOR CABLE TELEVISION DELIVERY SYSTEMS, incorporated herein by reference. Again, the channel changing is transparent to the viewer who believes the same channel is continuously being shown. The Multiple Channel method has the disadvantage of requiring that sufficient additional channels be available (by less bandwidth than Available Bandwidth method).

The Split Screen method transmits multiple commercials on a single channel using a split screen technique; commercials being pre-recorded and prepared prior to transmitting to the set top terminal 220. Although many commercials can be transmitted on a single channel, in the preferred form of the split screen method, only four commercials are shown. As the number of commercials increases the size and the amount of video information transmitted for each commercial decreases proportionately (i.e., 6, 8, 12, etc.). Using split screen methodology, either a masking technique or a scaling and repositioning of video technique must be used

at the set top terminal 220 to show the ad. The masking and repositioning-scaling techniques are further defined in co-pending application entitled, SET TOP TERMINAL FOR CABLE TELEVISION DELIVERY SYSTEMS, Ser. No. 08/160,193, owned by the assignee of the present invention and incorporated herein by reference. The scaling and repositioning technique produces better quality commercials, but requires expensive equipment at the set top terminal 220. The set top terminal 220 will perform audio switching with the split screen method to amplify the correct audio.

12. Alternatives to Basic Advertisement Targeting Routine

Figure 21 shows a software program flow 490 that is an alternative to the network controller's Basic Advertisement Targeting routine 374, depicted in Figure 17. The alternative program 490 allows each set top terminal 220 to be individually targeted with specific advertisements and is initiated automatically, block 492, by the network controller CPU 224 upon receipt of each polling response from a set top terminal 220. Thus, once the network controller 214 receives program access information from a set top terminal 220, the network controller CPU 224 begins the process of selecting a package of advertisements that is based on, among other things, that subscriber's demographic information and viewing history.

Upon receipt of a polling response from a set top terminal 220, the network controller CPU 224 reads the set top terminal identifier, 494, and the programs accessed, 496, from the polling response (or status report) (depicted in Figure 10b). The network controller 214 writes information on the programs accessed to the Program Scheduling database 320, updating the Access History File which contains listings of all programs accessed within the past week, month or year.

With continued reference to Figure 21, the network controller CPU 224 then calls a subroutine that sorts the programs accessed by program category, block 498. In turn, the program categories are sorted, 500, based on the number of times that programs appearing in each particular category are accessed. In so doing, this sorting subroutine determines and ranks those programs and program categories that are most frequently viewed by that set top terminal 220.

The subroutine can iteratively produce rankings for different time slots in a given day. In this way, different rankings can accommodate different viewing preferences during those time slots for a single set top terminal 220. For example, where rankings for eight three-hour time slots are desired, the subroutine determines a ranking of programs and program categories for each three-hour viewing period. Thus, a different ranking may be produced, for instance, for a morning time

slot and an evening time slot. All rankings of programs and program categories for that set top terminal 220 are written to the Viewer Profile database 314, updating the Viewer Log File, as at function block 502.

5 Next, the network controller CPU 224 calls a subroutine that correlates the updated Viewer Log File with the Advertisement Categories File in the Advertisement Library database 322, block 504. By
10 correlating these two files with one another, the subroutine assigns or correlates various categories of television commercials to each ranking of programs and program categories in the Viewer Log File. The
15 categories of television commercials and advertisements that may be so assigned are found in the Advertisement Categories File indicated generally at 354 as part of the
20 library 322 and may include: (1) Household Goods/Products, (2) Home Improvement and Maintenance, (3) Personal Hygiene, (4) Entertainment Items and Events, (5) Sporting Goods and Events, (6) Motor Vehicles and Related Products, (7) Foodstuffs and Beverages, and (8) Miscellaneous. Where, for example, the viewer has
25 watched a sporting event, the Sporting Goods and Events, Home Improvement and Maintenance, and Foodstuffs and Beverages categories may be assigned to that particular sporting event/program and Sports program category.

 Once the programs and program categories ranked in the Viewer Log File are correlated with the

advertisement categories in the Advertisement Categories File, the routine calls a sorting subroutine that ranks the groups of advertising categories correlated based on other information in the database files. In the preferred system, this ranking is primarily based on data in the updated Access History File and the updated Viewer Log File, as shown at function block 506. By using data on the viewer's past program selections and demographic information, the subroutine ranks the correlated categories of advertisements according to those likely to be of most interest to that viewer.

After the advertisement categories have been sorted and ranked, the routine selects the top three advertisement categories as the targeted categories for a given time slot and viewer, block 508. Individual advertisements are then chosen from the Advertisements File, with all selections made from the targeted categories, 510. The advertisements that are selected are written to the Advertisement Targeting File from where advertising packages can be generated, function 512, for transmission to the set top terminal 220. Such packages are generated by the network controller CPU 224, which accesses the Advertisement Targeting File and includes the targeted advertisements in the PCI signal. The entire routine is repeated for each set top terminal 220 and, alternatively, each viewer.

13. Account/Billing Routine

Figure 22 shows a software flow diagram for the network controller's Account/Billing routine 376, initiated automatically at block 520 by the network controller CPU 224 upon receipt of each polling response from a set top terminal 220. Upon receipt of such a response, the network controller CPU 224 identifies the set top terminal identifier from the polling response, block 522. The program access block in the polling response is also read, function 524, and the Access History File is updated with the received information, function 526. The routine then calls a subroutine that correlates the updated information in the Access History File with the Price Category File in the Program Library database, block 528. Once all programs accessed since the last polling cycle are assigned to a price category, the pricing information from each category is written to the Account History File, updating the file at 530. The network controller CPU 224 generates a billing report for each set top terminal 220 based on the updated account history, function 532. This billing report can be sent to the set top terminals 220 in a polling request. Specifically, in one embodiment, the information field of the frame format described in Figure 9a is used to provide the set top terminal 220 with billing information.

Account information for each set top terminal 220 can be viewed through a monthly account review menu.

The account information necessary to create the monthly account review revenue may be stored either in the memory of the set top terminal 220 or at a remote location that communicates with the set top terminal 220. In the simplest embodiment, the set top terminal 220 records a subscriber's selections locally and calculates the monthly account review based upon the subscriber's selections which require the payment of fees. This monthly account information is stored locally and sent to the network controller 214 upon polling.

The Account/Billing routine is capable of processing account and billing information generated in other embodiments. For example, in an alternate embodiment, the subscriber's viewing selections and billing information may be continuously maintained at the network controller 214 or a remote site connected via communication lines to the cable headend 208. The network controller 214 or the remote site must regularly transmit the monthly account information to the set top terminal 220.

Each embodiment, such as local billing storage at the set top terminal 220, billing by the network controller 214 or billing by a remote site, has advantages and disadvantages. If the account information and processing is done locally at the set top terminal 220, each set top terminal 220 must be provided with the memory and necessary processing capability to maintain

the account. This greatly increases the cost of a set top terminal 220. If the account information is maintained remotely, the remote site must remain in regular contact with the set top terminal 220 in order to provide the subscriber with billing information. To accommodate homes with multiple viewers two or more set top terminals 220 may be placed on a single bill or two accounts may be created for one set top terminal 220.

Figure 23 shows another embodiment in which billing may be accomplished through the use of remote statistical and billing sites (SBS). In this arrangement, statistical and billing information from individual communities of set top terminals 1750 is communicated through cable headend sites to regional statistical and billing sites 1730 (SBS). A regional SBS may serve several cable headend sites, shown at 1732. The regional SBS 1730 calculates billing and statistical information and passes necessary billing information back downstream through the network controller 214 at the cable headend 208 to an appropriate single set top terminal 220 in a subscriber's home. In addition, the regional SBS 1730 communicates the billing and statistical information received on program viewer choices to the central SBS 1740.

The central SBS 1740 accumulates the data received from a number of regional statistical and billing sites and calculates national statistical and billing

information. In the preferred embodiment, the regional SBS 1730 prints and mails bills to subscribers. The central SBS 1740 can calculate program ratings, shares and HUTS (homes using televisions) for the nation and by region. With information from interactive TV programs, sophisticated statistical information may be gathered through the network controllers of the cable headends.

This arrangement for billing and statistical information provides the operators of the system with the advantages of distributive processing.

Remote billing sites may serve regions of the country by having each cable headend 208 in a region of the country connected to one regional billing site. The information from the regional billing sites may then be communicated on a less frequent basis to the operations center 202 or a central billing location. This method of distributed processing of billing enables the central billing location to receive fewer communications and be more efficient. In addition, the communication links between the cable headend's network controller 214 and regional sites will be of shorter distance than communication links to the operations center 202 from the cable headends 208. This should result in a cost savings to the system operator.

Regional statistical and billing may, however, be eliminated and all communications from the cable headend 208 may proceed to the Central SBS 1740. In

fact, the Central SBS 1740 can be collocated with the operations center 202 and all functions performed at one central location. If the cable program packaging and delivery system 202 is established in just one locale, the network controller 214 can perform all the statistical and billing procedures.

The terms and descriptions used herein are set forth by way of illustration only and are not meant as limitations. Those skilled in the art will recognize that numerous variations are possible within the spirit and scope of the invention as defined in the following claims.



THE CLAIMS DEFINING THE INVENTION ARE AS FOLLOWS:

1. A method for targeting advertising to at least one subscriber, the method including the steps of:

gathering programs watched data from a subscriber;

analyzing the programs watched data to determine the frequency of programs watched by the subscriber;

correlating the analyzed programs watched data with categories of advertisements, wherein each advertisement category includes at least one advertisement;

selecting an advertisement from the correlated advertising categories;

and

transmitting the selected advertisement for display to the subscriber.

2. An apparatus for data processing related to program advertisements wherein data on viewer profiles is used, including:

a receiver for receiving data, wherein viewer profile data is received from one or more remote locations;

a first memory location, operably connected to the receiver, wherein information on advertisements is stored;

a second memory location, operably connected to the receiver, wherein data on viewer profiles is stored; and

a central processing unit, operably connected to the first and second memory locations, wherein the information on advertisements and viewer profile data are processed, and at least one advertisement is scheduled, whereby at least one advertisement is scheduled to be inserted into programming.

3. The apparatus of claim 2 wherein programs watched data is used, further including:

a third memory location, operably connected to the central processing unit, wherein program watched data is stored; and wherein the central processing unit processes the information on advertisements, viewer profile



data and program watched data to schedule advertisements.


4. A system for targeting advertising to at least one subscriber including:
 a means for gathering programs watched data from a subscriber;
 a processor, operably connected to the gathering means, including
 a means for analyzing gathered programs watched data to determine
 the frequency of programs watched by the subscriber;
 a means for correlating the analyzed programs watched data with
 categories of advertisements, wherein each advertisement category includes at
 least one advertisement; and
 a means for selecting an advertisement from the correlated advertising
 categories; and
 a transmitter, wherein the selected advertisement is transmitted.

5. The system of claim 4 further including a display, wherein the transmitted
 advertisement is displayed.

6. The system of claim 4 wherein the means for analyzing counts the
 programs watched data to determine the frequency of programs watched by the
 subscribers to be targeted, and wherein the counts are arranged in at least one
 programs watched matrix.

7. The system of claim 3, wherein the programs watched matrix is arranged
 by program category and time slot and the analyzing means includes:
 a means for reading the programs watched matrix; and
 a means for sorting the programs watched counts within a time slot from
 highest to lowest.

8. The system of claim 5 wherein the processor is a network controller and
 the display is a set top terminal operably connected to a television.



9. The system of claim 8 wherein the transmitter transmits the selected advertisements on a television channel, and the set top terminal is capable of changing television channels.

10. The system of claim 8 wherein the selecting means selects a plurality of advertisements, and the transmitter transmits the selected advertisements on a single television channel, and the set top terminal is capable of masking undesired video.

11. The system of claim 8 wherein the selecting means selects a plurality of advertisements, and the transmitter transmits the selected advertisements on a single television channel, and the set top terminal scales and repositions video.

12. The system of claim 4 wherein the processor polls the set top terminals to monitor the current program being watched by the subscriber for the occurrence of commercial breaks, further including:

a means for polling the set top terminals; and

wherein the transmitter transmits the selected advertisements in response to the occurrence of commercial breaks determined by the polling means.

13. The system of claim 4 wherein the means for gathering programs watched data is a set top terminal including:

memory wherein programs watched data is stored; and

the processor further includes a means for receiving the stored programs watched data from the set top terminal.

14. The system of claim 13, wherein the processor polls the set top terminal and wherein the receiving means receives status reports sent by the set top terminal and the receiving means further includes:

a means for storing the programs watched data;

a means for generating a polling request message that directs the set top terminal to initiate transmission of set top terminal status reports containing



programs watched information;

a means for report processing, wherein the received set top terminal status reports are processed to produce polling response data; and

a means for temporarily storing the polling response data.

15. The system of claim 14, wherein the means for report processing includes:

a means for demodulating the received set top terminal status reports, whereby demodulated set top data is produced;

a means for reading at least one information field in the demodulated set top data, wherein each information field is appended to a set top terminal identification number field;

a means for sorting each information field in the received set top terminal status reports by set top terminal identification number; and

a temporary memory means, wherein the sorted information fields for each set top terminal is accumulated, and wherein the accumulated sorted information fields produce the polling response data.

16. The system of claim 15 further including a means for updating the programs watched data in the storing means with the polling response data.


17. The system of claim 4 further including:

a means for gathering demographic data from subscribers to be targeted;

and

wherein the means for correlating also correlates the demographic data with the categories of advertisements.

18. The system of claim 17 wherein there is a group of subscribers that forms a statistically significant number of subscribers, and the means for gathering demographic data uses both gathered demographic data and simulated demographic profiles as the demographic data and the means for gathering demographic data further includes:



a means for gathering demographic data from the statistically significant group of subscribers; and

a means for generating a simulated demographic profile of the subscribers to be simulated by comparing the programs watched data of the subscribers to be simulated with the gathered demographic data and the programs watched data of the statistically significant group of subscribers.

19. A system for targeting advertising including:

a means for gathering programs watched data from subscribers to be targeted;

a processor including:

a means for counting the gathered programs watched data to determine the frequency of programs watched by the subscribers to be targeted, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot; and

a means for creating subscriber group information indicating a group assignment for each subscriber to be targeted by correlating the programs watched counts with the categories of advertisements, wherein each advertisement category includes advertisements available for targeting to the set top terminals;

a transmitter, wherein the subscriber group information is transmitted to the subscriber's set top terminal in a control information stream that instructs the set top terminal in selecting targeted advertisements for display during viewing of programs; and

a set top terminal capable of selecting targeted advertisements for display.

20. The system of claim 19 wherein the processor is a network controller.

21. The method of claim 1 wherein at least one advertisement within an advertisement category is a promotion.



22. The method of claim 1 wherein at least one advertisement within an advertisement category is an infomercial.

23. The apparatus of claim 2 wherein the information on advertisements identifies an infomercial.

24. The apparatus of claim 2 wherein the information on advertisements identifies a promotion.

25. The apparatus of claim 2 wherein the information on advertisements identifies the individual advertisements by name.

26. The apparatus of claim 2 wherein the information on advertisements identifies the individual advertisements by type of advertisement.

27. The apparatus of claim 2 wherein the information on advertisements includes directions for inserting advertisements into a program line-up.

28. The apparatus of claim 2 further including:
a third memory location, operably connected to the processor and receiver, wherein the third memory location stores video corresponding the information on advertisements; and

wherein the central processing unit processes the video with the information on advertisements and viewer profile data.

29. The apparatus of claim 2 wherein the apparatus is located at an operations center.

30. The apparatus of claim 2 wherein the apparatus is located at a cable headend.



31. The apparatus of claim 30 wherein the first memory location also stores video corresponding to the information on advertisements; and further including:

a second receiver, wherein the second receiver receives television programming and insert directions, and whereby the insert directions are used by the cable headend to insert an advertisement into the television programming, and the central processing unit schedules at least one advertisement based on the insert directions.

32. The system of claim 4 wherein at least one advertisement within an advertisement category is an infomercial.

33. The system of claim 4 wherein at least one advertisement within an advertisement category is a promotion.

34. The system of claim 4 wherein the system is located at a cable headend.

35. The system of claim 4 wherein the system is located at an operations center.

36. The system of claim 4 wherein the gathering means gathers programs watched data from a plurality of subscribers, wherein the processor processes the programs watched data gathered from a plurality of subscribers; and wherein the transmitter transmits the selected advertisements to the plurality of subscribers.


37. The system of claim 5 wherein the display displays a menu of the selected advertisements from which the subscriber chooses an advertisement to watch.

38. The system of claim 5, further including:

a receiver, operably connected to the display, wherein the receiver receives television signals.



39. The system of claim 4 wherein the processor further including:
a memory, operably connected to the receiver, wherein the memory stores video corresponding to the selected advertisement.
40. A method of targeting advertisements including the steps of:
gathering programs watched data from a subscriber;
analyzing gathered programs watched data to determine the frequency of programs watched by the subscriber;
correlating the analyzed programs watched data with a category of advertisements, where each advertisement category identifies at least one advertisement; and
selecting an advertisement from an advertisement category based on programs watched data.
41. The method of claim 40 further including the step of storing the advertisement identified by the advertising category.
42. The method of claim 40 wherein at least one advertisement within an advertisement category is a promotion.
43. The method of claim 40 wherein at least one advertisement within an advertisement category is an infomercial.
44. The method of claim 40 further including the steps of:
receiving several channels of advertisements by broadcast, wherein each channel corresponds to an advertisement category and each advertisement category identifies at least one advertisement;
wherein the selected advertisement is within a channel; and
switching to the channel containing the selected advertisement.



45. The method of claim 44, wherein the switching is transparent to the viewer.

46. An apparatus for targeting advertising to a subscriber, the apparatus including:

means for gathering programs watched data from a subscriber;

means, operably connected to the gathering means, for analyzing the programs watched data;

means, for correlating the analyzed programs watched data with information on advertisements, wherein the information on advertisements identifies at least one advertisement; and

means for selecting an advertisement based on the correlation.

47. The apparatus of claim 46 wherein the information on advertisements identifies an infomercial.

48. The apparatus claim 46 wherein the information on advertisements identifies a promotion.

49. The apparatus of claim 46 wherein the apparatus is a set top terminal.

50. The apparatus of claim 49, further including a receiver, wherein the receiver receives television signals.

51. The apparatus of claim 50 further including:

a memory, wherein the memory stores video corresponding to the selected advertisement; and

wherein the receiver also receives insert directions which instruct the set top terminal to insert the corresponding video.

52. The apparatus of claim 50 wherein the receiver receives several channels of advertisements, and the apparatus further including:



means for selecting a channel of advertisements;
 means for switching to the selected channel; and
 a display, wherein the switching of channels is transparent to the viewer.

53. The apparatus of claim 46 further including:
 means, operably connected to the correlating means, for storing the
 information on advertisements.

54. The apparatus of claim 53, wherein the storing means stores video
 corresponding to the selected advertisement.

55. The apparatus of claim 46 wherein the gathering means, analyzing
 means, and correlating means are located at an operations center; and wherein
 the selecting means is located at a cable headend.

56. The apparatus of claim 55 wherein the correlating means generates an
 output further including:

a transmitter, operably connected to the correlating means, wherein the
 transmitter transmits the output;

a receiver, operably connected to the selecting means, wherein the
 receiver receives the output; and


a memory, operably connected to the selecting means, wherein the
 memory stores video of the selected advertisement.

57. An apparatus for creating a program line-up targeted to a subscriber,
 wherein data on programs watched is used to determine the subscriber's
 preferences, including:

means for gathering programs watched data from the subscriber to be
 targeted;

means, operably connected to the gathering means, for analyzing
 programs watched data;

means for developing a program line-up based on the analyzed data;



and

means, operably connected to the developing means, for transmitting the program line-up to the subscriber.

58. The apparatus of claim 57 further including:

means for gathering marketing data; and wherein the analyzing means also analyzes the marketing data.

59. A method of targeting programs including the steps of:

gathering programs watched data from a subscriber;
analyzing the gathered programs watched data to determine viewing habits of the subscriber;
storing at least one program for selection at a remote location;
selecting at least one program based on the analyzed data; and
transmitting at least one selected program from the remote location to the subscriber.

60. A cable headend for targeting programs including:


means for gathering programs watched data from a subscriber;
means, operably connected to the gathering means, for analyzing the gathered programs watched data to determine viewing habits of the subscriber;
means for selecting a program based on the analyzed data;
means, operably connected to the selecting means, for storing programs, wherein the selected program is stored.

61. The apparatus of claim 60 further including:

means, operably connected to the selecting means, for displaying the selected program to the subscriber.

62. An apparatus for targeting video, wherein data on programs watched is used to determine a subscriber's preference, including:

means for gathering programs watched data from a subscriber;



means, operably connected to the gathering means, for analyzing the gathered programs watched data to determine a subscriber's preference;

means for selecting video based on the analyzed data;

a memory, operably connected to the selection means, wherein video is stored; and

means, operably connected to the memory for transmitting the selected video to the subscriber.

63. The apparatus of claim 52 wherein the video is television programming.

64. The apparatus of claim 52 wherein the video is advertisements.

65. The method of claim 1, further including displaying the selected advertisement.

66. The method of claim 1, wherein analyzing programs watched data includes counting programs watched by the subscriber and further including arranging the counts in at least one programs watched matrix.

67. The method of claim 66, wherein the programs watched matrix is arranged by a program category and a time slot and wherein the counts are arranged within the time slot from highest to lowest.

68. The method of claim 65, wherein the selected advertisement is displayed on a display of a television.

69. The method of claim 1, wherein the transmitted advertisements are transmitted on a television channel.

70. The method of claim 1, wherein selecting the advertisement selects a plurality of advertisements, the plurality of advertisements being transmitted on a single television channel, and further including masking undesired



advertisements out of the plurality of advertisements.

71. The method of claim 1, wherein selecting the advertisements selects a plurality of advertisements, the plurality of advertisements being transmitted on a single television channel, a desired advertisement of the plurality of advertisements being scaled and repositioned on a display.

72. The method of claim 1, wherein the advertisements are transmitted in response to an occurrence of commercial breaks within a program.

73. The method of claim 1, wherein the transmitted advertisements are stored in a set top, the advertisements being displayed based on an occurrence of commercial breaks within a program.

74. The method of claim 1, wherein the programs watched data is gathered by a set top, the set top storing the programs watched data.

75. The method of claim 74, wherein the set top analyzes the programs watched data.

76. The method of claim 74, wherein a processor coupled to the set top receives the programs watched data and analyzes the programs watched data.

77. The method of claim of claim 76, further including:
 generating a polling request message that directs the set top to transmit a set top status report that contains the programs watched data;
 processing the received set top status reports to produce polling response data; and
 storing the polling response data.

The method of claim 77, wherein processing the received status reports includes:

78
79
80
81
82
83
84
85
86
87
88
89
90
91
92
93
94
95
96
97
98
99
100



demodulating the received status reports, producing demodulated set top data;

reading at least one information field in the demodulated set top data, each information field being appended to a set top identification number field; and

sorting each information field by set top identification number.

79. The method of claim 78, further including updating the programs watched data based on the processed status reports.

80. The method of claim 1, further including:
gathering demographic data related to the subscribers; and
correlating the demographic data with the categories of advertisements.

81. The method of claim 80, wherein gathering the demographic data includes one of sending questionnaires to the subscribers, receiving demographic data via the set top and gathering demographic data from third parties.

82. The method of claim 80, wherein gathering demographic data includes generating a simulated demographic profile of the subscribers by comparing the programs watched data of the subscribers to a sample viewer profile, the sample viewer profile constructed based on demographic data received from a statistically significant sample of subscribers.

83. The method of claim 1, wherein a display displays a menu of the selected advertisements from which a subscriber chooses an advertisement to watch.

84. The method of claim 76, wherein the set top transmits a set top status report that contains the programs watched data to the processor using a random access method.



85. The method of claim 84, wherein the random access method further includes a CDMA/CD protocol.

86. The method of claim 76, wherein the set top transmits a set top status report that contains the programs watched data to the processor using a standard telephone line.

DATED this 18th day of November, 1999

DISCOVERY COMMUNICATIONS INC


WATERMARK PATENT & TRADEMARK ATTORNEYS
290 BURWOOD ROAD
HAWTHORN VICTORIA 3122
AUSTRALIA

RCS:SMM:SLB DOC 21 AU6086798.WPC

8
6
7
9



Fig. 3



6 04 00 00007




Fig. 4

6 0 4 3 0 0 0 5 7




Fig. 5

6 0 4 2 5 8 4




Fig. 6b




Fig. 7

Fig. 8a




Fig. 8b





Fig. 8c



5 14 30 SUBS




Fig. 9a

5 0 4 3 0 0 0 7

13/27




Fig. 9b

0 4 9 0 0 0 0 7




Fig. 11




Fig. 12

5
6
7
8
9




Fig. 13




Fig. 14

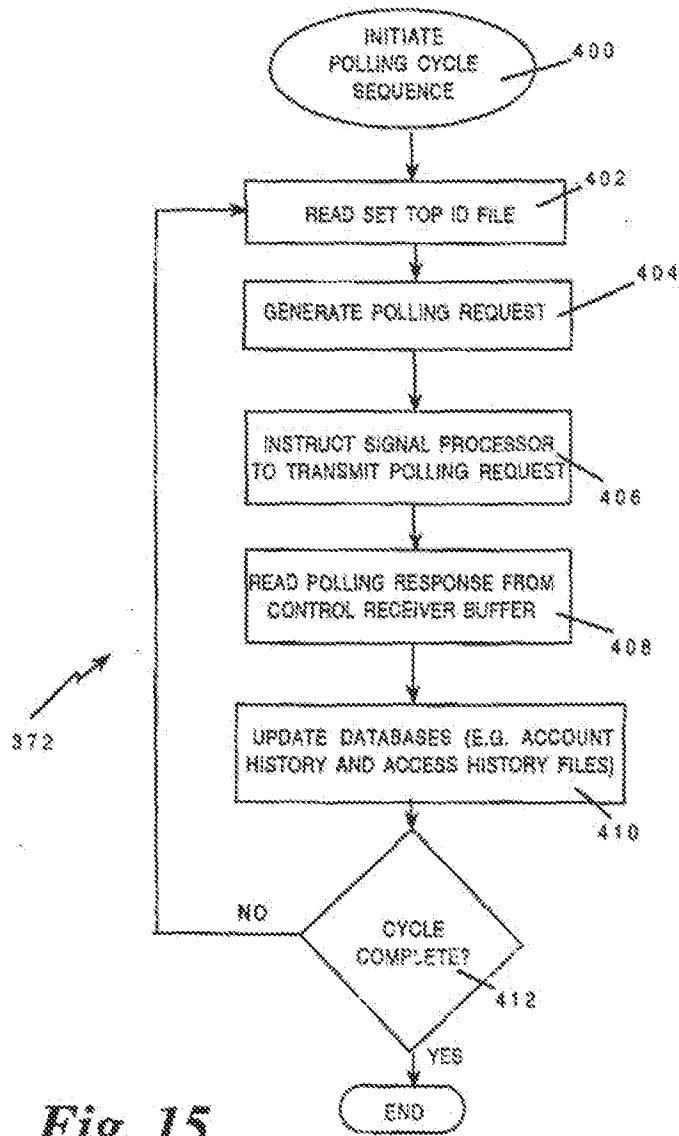


Fig. 15

FIG. 16

20/27

351

TIME SLOT	SPORTS	NEWS	MOVIES	CHILDREN	ENTERTAINMENT
0000-0400	0	2	2	0	0
0400-0800	0	2	0	0	0
0800-1200	2	0	10	2	8
1200-1600	20	1	9	5	4
1600-2200	8	8	13	0	5
2200-2400	0	10	2	4	2

Fig. 16



374




Fig. 17

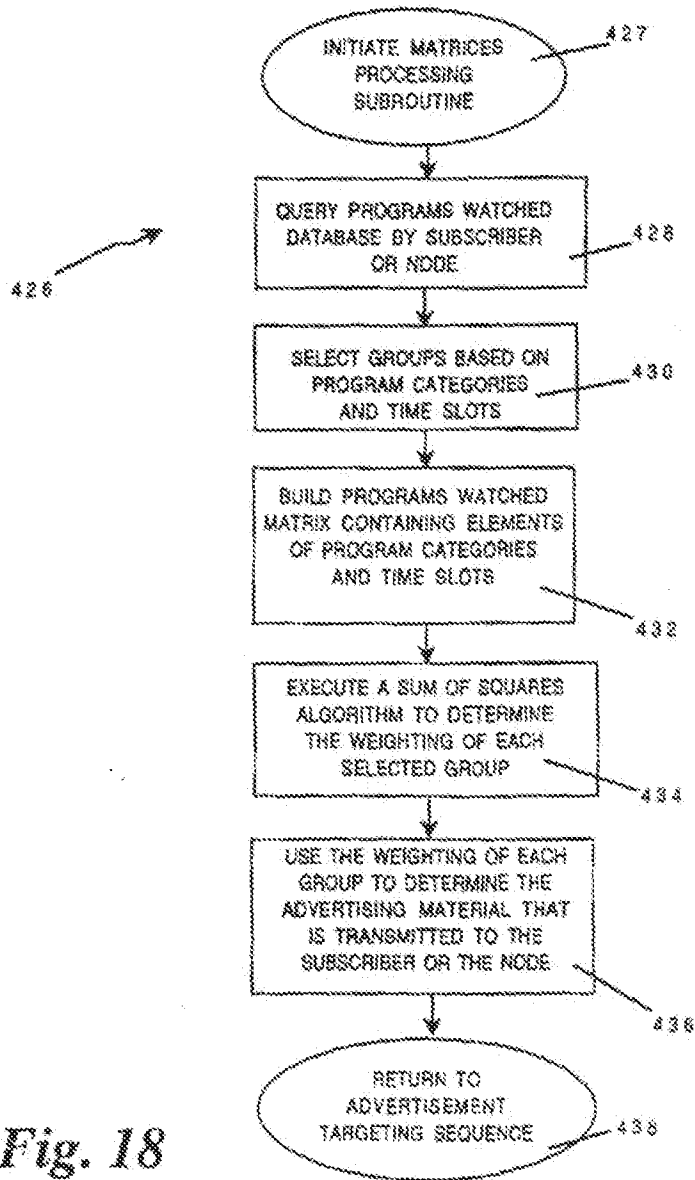


Fig. 18




FIG. 200




Fig. 200a

Fig. 200b

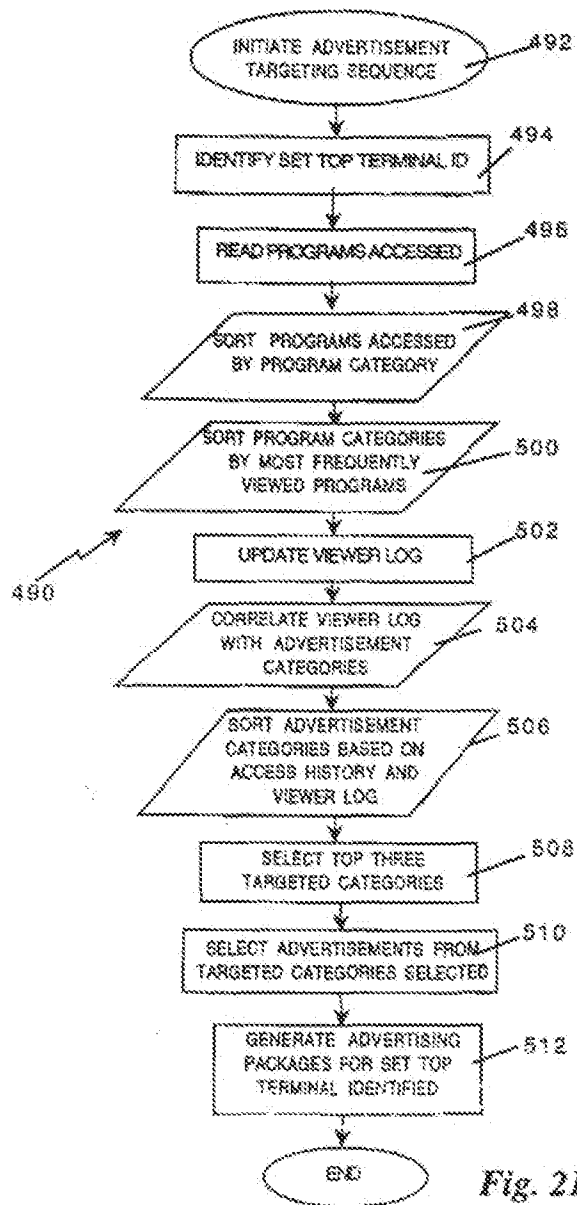


Fig. 21

376






Fig. 22

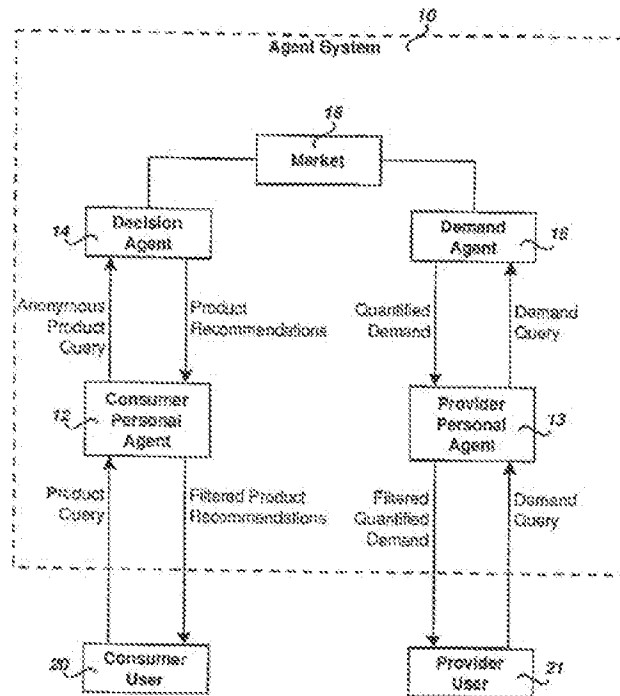
6 0 4 3 0 0 0 0 0

21/27





(72) PECKOVER, Douglas L., US
(71) PERSONAL AGENTS, INC., US
(51) Int.Cl.⁵ G06F 17/60
(30) 1996/01/17 (60/010,087) US
(54) AGENTS INTELLIGENTS POUR COMMERCE
ELECTRONIQUE
(54) INTELLIGENT AGENTS FOR ELECTRONIC COMMERCE



(57) La présente invention concerne un système de commerce électronique (10) comportant des agents personnels (12, 13) représentant des consommateurs et des fournisseurs sur un marché virtuel (28). Les agents personnels consommateurs cachent l'identité du consommateur et sont capables de créer des agents de décision (14) achetant des produits et aidant les consommateurs à comparer et classer les produits. Les agents personnels fournisseurs sont capables de créer des agents de demande (16) qui quantifient la demande et des

(57) A system for electronic commerce (10) having personal agents (12 and 13) that represent consumers and providers in a virtual marketplace (28). Consumer personal agents conceal the identity of the consumer and are capable of creating decision agents (14) that shop for products and assist consumers in comparing and ranking products. Provider personal agents are capable of creating demand agents (16) that quantify demand and target specific consumers without learning the identity of the consumers. Based on data generated by the activities





(2D) (A) **2,242,874**
(86) 1997/01/17
(87) 1997/07/24

consommateurs cibles spécifiques sans avoir à apprendre l'identité des consommateurs. Partant des données générées par les activités des agents de décision et des données de préférence tenues à jour par les agents personnels consommateurs, les agents personnels fournisseurs peuvent quantifier les demandes présente, passée et future, simuler la demande, et peuvent quantifier les consommateurs cibles spécifiques auxquels adresser la publicité et d'autres messages. Ces agents personnels fournisseurs peuvent coopérer avec les agents personnels consommateurs pour saisir des données sur les raisons motivant les ventes effectuées ou manquées et proposer aux consommateurs des paiements à contreparties. Les agents personnels consommateurs peuvent refuser automatiquement les messages non sollicités qui ne satisfont pas à leurs préférences de consommateurs.

of the decision agents and on preference data maintained by consumer personal agents, provider personal agents can quantify current, historical, and future demand, simulate demand, and target specific consumers for advertising and other messages. Provider personal agents can cooperate with consumer personal agents to collect data about reasons for sales and lost sales and to offer consideration payments to consumers. Consumer personal agents can automatically reject unsolicited messages that do not satisfy the consumer's preferences.





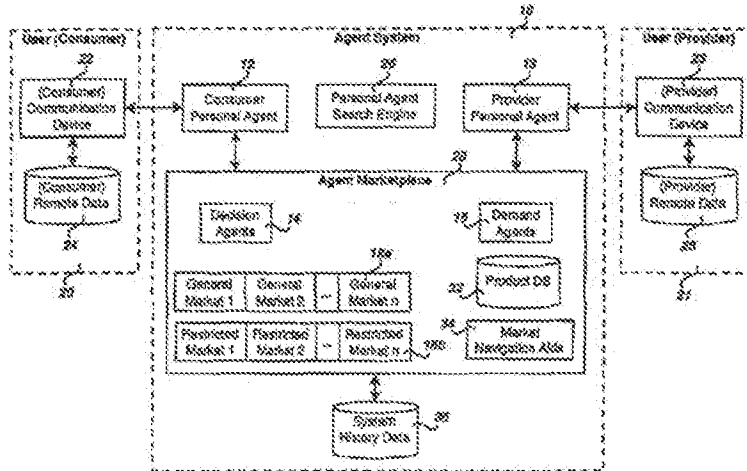
PCT

WORLD INTELLECTUAL PROPERTY ORGANIZATION
International Bureau

INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(51) International Patent Classification ⁶ : G06F 17/60, G06G 7/52		A1	(11) International Publication Number: WO 97/26612
			(43) International Publication Date: 24 July 1997 (24.07.97)
(21) International Application Number: PCT/US97/01057		(81) Designated States: AU, CA, CN, IL, JP, KR, MX, Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, CH, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).	
(22) International Filing Date: 17 January 1997 (17.01.97)			
(30) Priority Data: 60/010,087 17 January 1996 (17.01.96) US Not furnished 30 December 1996 (30.12.96) US		Published With international search report.	
(71) Applicant: PERSONAL AGENTS, INC. [US/US]; Suite 313, 14651 Dallas Parkway, Dallas, TX 75240 (US).			
(72) Inventor: PECKOVER, Douglas, L.; Personal Agents, Inc., Suite 313, 14651 Dallas Parkway, Dallas, TX 75240 (US).			
(74) Agent: TIMMONS & KELLY; The White House on Turtle Creek, 2401 Turtle Creek Boulevard, Dallas, TX 75219-4760 (US).			

(54) Title: INTELLIGENT AGENTS FOR ELECTRONIC COMMERCE



(57) Abstract

A system for electronic commerce (10) having personal agents (12 and 13) that represent consumers and providers in a virtual marketplace (20). Consumer personal agents conceal the identity of the consumer and are capable of creating decision agents (14) that shop for products and assist consumers in comparing and ranking products. Provider personal agents are capable of creating demand agents (15) that quantify demand and target specific consumers without learning the identity of the consumers. Based on data generated by the activities of the decision agents and on preference data maintained by consumer personal agents, provider personal agents can quantify current, historical, and future demand, simulate demand, and target specific consumers for advertising and other messages. Provider personal agents can cooperate with consumer personal agents to collect data about reasons for sales and lost sales and to offer consideration payments to consumers. Consumer personal agents can automatically reject unsolicited messages that do not satisfy the consumer's preferences.

-1-

DescriptionIntelligent Agents for Electronic Commerce5 Background of the Invention1. Cross-References to Related Applications

This application is related to Provisional Patent Application Serial Number 60/010,087, Filed 17 January 1996. This application is also related to Provisional Patent Application Serial Number [XXXXXXXXXX], Filed 30 December 1996.

10

2. Field of the Invention

The present invention relates to the gathering and analysis of market transaction data, where such transactions are contemplated or completed by electronic means, and specifically to the use of software agents to represent and to assist the activities of consumers and providers within an electronic "virtual marketplace".

15

3. Description of Related Art

The trading of goods and services is one of the basic human activities. From the first meeting of pre-historic men to trade animal skins for berries, to country fairs, to the most intricate computer-based trading of international financial instruments, the marketplace supports the survival and flourishing of our lives.

20

Every era has established a marketplace based on the technology of its time. Early marketplaces were locations for face-to-face meetings between people offering items for trade. The development of printed material and postal correspondence made it possible for people to trade without ever meeting in

25

-2-

person. In modern times, a sophisticated division of labor system involving producers, suppliers, distributors, advertisers, market researchers, payment clearinghouses, customers, and many others, supports our economic lives.

The flow of goods and services, however, is not the whole story. The people and organizations on the production, distribution, sales, and marketing side of trade ("providers") need to understand what goods and services are desired by buyers. The people and organizations who purchase or acquire offered items ("consumers") need to learn what goods and services are available. The flow of this market information is critical to the successful operation of a market.

A brief consideration of modern methods of commerce shows us that consumers have many ways to gather information about available goods and services. Some of these are: visiting stores; browsing catalogs; viewing advertisements on television, on billboards, and in magazines and newspapers; soliciting recommendations from friends; and receiving unsolicited mail advertisements.

We also observe that vendors and other providers use many avenues to disseminate information about available goods and services. These include: broadcast and direct advertisements; in-store displays; telephone solicitations; and so on.

Providers also try to collect information about consumers' desires and buying habits so that they can better serve their current customers and gain new ones, thereby enhancing profits. Current market research methods include: keeping current customer purchase data on file; buying or renting mailing lists of other vendors; listening to focus groups; running pilot sales in test markets; and so on.

However, even with so many pathways for information exchange, there are many disadvantages inherent in the current methods of commerce.

-3-

Disadvantages of Current Methods of Commerce

Although consumers clearly reap many benefits in today's marketplace, there are still many disadvantages that lead to consumer frustration. Here are some obvious disadvantages for consumers:

- 5 * Just collecting basic information about available goods and services, their features, and their prices is tedious and time-consuming. To find a particular product or verify availability, price, and features, the consumer typically may need to visit several stores across a wide area. Another consumer may choose to spend hours on the telephone, perhaps incurring various communication charges,
- 10 navigating irritating automated attendants, trying to reach a human who can answer questions. Yet another consumer will peruse stacks of catalogs and accumulated direct mail advertisements. This process wastes the consumer's valuable time, especially when the search fails.
- * Information isn't always available when it's needed. Consumers
- 15 frequently rely on ephemeral television and radio advertisements to learn of products and special promotions. But these ads seldom arrive when the consumer is ready to make a selection. Even with print ads, the information is soon lost as the stack of newspapers is carried away for recycling.
- * It is very difficult for consumers to pick out items of interest from
- 20 the daily bombardment of advertising. Television and radio spots, billboards, street corner kiosks, the daily newspaper, direct mail coupon packages, in-store merchandise displays, magazine advertisements and inserts, posters on and in mass transportation vehicles, storefront banners - these advertisements and a host of others vie for the consumer's attention daily. Marketing research firms use the
- 25 term "impression" to denote one occurrence of a person perceiving an advertisement. Some authorities claim that an urban consumer in New York City receives up to 15,000 impressions per day! It is no wonder that consumers

-4-

become desensitized to the advertising around them, and fail to see the items that are truly of interest.

* Consumers have little or no control over the timing or presentation of advertising. They can't arrange to receive information at a convenient time.

5 Telephone solicitors always seem to interrupt dinner. The radio advertisement for an anticipated concert is on the air during "drive time" - when the consumer is driving and can't write down the phone number of the ticket agency. Consumers can't even choose a preferred format for receiving advertisements. If the merchant advertises only in broadcast media, the consumer won't be able to find the

10 information in today's newspaper. Sorting through the barrage of direct mail may require more time than the consumer can allocate to this task, and so the consumer discards a potentially useful notice.

* A consumer usually receives no direct benefit for inspecting a provider's advertisements, except for the information itself if the advertisement is

15 relevant. Occasionally, providers and marketing research firms give consumers a small fee or gift in return for participating in a survey. This is called "paying a consideration." However, there are few opportunities for consumers to receive considerations, and no way for consumers to seek out providers that are willing to pay considerations.

* When a consumer has an immediate need for product or service

20 information, it may be nearly impossible to gather the information quickly. This is especially true when the consumer doesn't know where to look for the information, or doesn't have quick and convenient access to sources of information.

* Many consumers rely on the recommendations or evaluations of

25 third parties to help them make buying decisions. Consumer rating guides and

-5-

endorsers provide this kind of information, but the information is not always available at the time or place of the consumer's purchase decision.

Providers, too, have their share of frustrating difficulties using the current methods of commerce. Some of the disadvantages for providers are:

5 * Much of the demographic data from traditional sources is out-of-date or incorrect. Providers frequently buy or rent magazine subscription lists and other providers' customer lists. However, a name on a list is not necessarily removed when the individual no longer subscribes or is no longer a customer. It is expensive to examine lists for data entry errors, duplicates, minor
10 spelling variations, and so on. When a provider exchanges a list with another provider, who exchanges with a third provider, who exchanges with yet another provider, a provider can easily end up with a list that contains a high proportion of names of consumers who are not at all interested in the provider's goods.

 * It is expensive to target advertisements to specific customers. As
15 mentioned above, mailing lists frequently contain a high percentage of names of uninterested customers, yet buying or renting mailing lists is costly. Controlled circulation magazines, where potential subscribers complete a qualification form to receive a (usually complimentary) subscription, don't always yield truly qualified names, since there typically is no mechanism to verify the information that the
20 potential subscribers report.

 * It is difficult for a provider to deliver information directly to a consumer who is ready to buy. Even if providers could easily identify those consumers, there is a time lag for the delivery of pertinent information. For example, direct mail requires planning weeks in advance. An in-store salesperson
25 can assist the consumer, but only if the consumer has previously learned that the store carries suitable products. Providers need a mechanism for delivering information to consumers precisely at the moment when it is most helpful.

-6-

* It is difficult and costly to personalize information to a consumer. Most consumers are unimpressed by feeble attempts at personalization such as form letters that read "Dear MR. JONES JOHN, The JOHN FAMILY may have already won..."

5 * It is especially difficult to determine what specific interests a particular consumer has. Most data of this nature must be inferred from subscription lists, member lists, spotty purchase history, etc. It is rare that a consumer directly informs a provider of a particular interest.

10 * Providers have little control over the timing of the delivery of their printed advertising messages. The use of special mailing classes for bulk mail to reduce mailing costs results in erratic delivery times. In the United States, providers using Third Class mail cannot pinpoint even the week that the mail will be delivered. For example, sometimes consumers don't receive a special sale notice until after the sale date. Also, mailed notices must be prepared well in
15 advance to take advantage of bulk mail, so the provider's quick response to market conditions is impeded.

* It is difficult for providers to verify delivery of their messages. Did the direct mail piece get read by the consumer? Was it even delivered? Did interested consumers view the television advertisement? How many consumers
20 noticed the billboard? How many consumers read the newspaper or magazine notice? Market survey and research firms attempt to measure delivery, but their methods are necessarily statistical since they can't survey every household in the target area. Even with their limited usefulness and contested accuracy, these services are expensive.

25 * The low success rates of direct mail (typically 2% to 4%) wastes much of the natural resources for printing and distributing the mailing, as well as wasting the money for preparation of the direct mail. Huge amounts of unsolicited

-7-

mail end up in the wastebasket. Some authorities estimate that up to 70% of unsolicited mail is never opened.

* It is expensive and cumbersome for providers to offer "considerations" to consumers. A consideration is a payment or award of some value to a consumer in return for a consumer viewing an advertisement or participating in a marketing survey. Sometimes considerations are given to encourage the consumer to an action, as when marketing research firms include a dollar bill in an unsolicited direct mail survey. Not only is there the cost of the consideration itself, there are the additional costs to the provider of identifying potential consumer recipients and preparing some means of delivery such as direct mail.

* Providers have no practical way to get real-time (immediate) feedback on the success of their promotions. Marketing research on a particular product typically requires at least several weeks or months and is very costly. This leads marketers to test only large product groups and discourages them from gathering data about individual products.

* Providers typically have no way to collect information about why a consumer purchased an item. Providers run special promotions, and consumers buy things, but it is arduous, tedious, and error-prone to draw the connection. It is difficult to judge the effects of promotions targeted to different sets of consumers, because providers don't know which promotion persuaded the consumer to buy.

* Providers have few ways to collect information about "lost sales", or why a consumer did not purchase an item. If lost sale data could be determined, a provider could better tailor offers to the consumer's needs and desires.

* Providers attempt to measure and predict "consumer demand" to help determine the number and mix of products and services to offer and the prices

-8-

to charge. Consumer demand is a measure of the number of consumers who want to purchase a product or service. It is typically calculated on a large scale with statistical models using historical purchase data. Demand can only be calculated based on purchases that have already occurred, since providers have few
5 mechanisms to determine what new items consumers might want, or at what price consumers would buy. Demand information would be much more useful if it could guide providers into new territory, or if it could warn providers that a planned product would likely have few buyers or is being offered at an unsuitable price. With more accurate demand information, providers could plan inventories that are
10 better matched to consumer desires, resulting in fewer markdowns.

- Providers have no mechanism for using actual demand data to simulate consumer demand under varying conditions. It would be useful if providers could run "what if" scenarios to see the effects on demand of different prices, varying packaging, special discounts, etc. Using actual historical and
15 current data to calculate this "theoretical demand" would be more accurate than using estimated or aggregate data.

The Promise of Electronic Commerce

The unexpected and explosive growth and popularity of the Internet in
20 recent years has opened a new avenue for commerce - "electronic commerce". Electronic mail ("e-mail"), the delivery of messages via electronic communication networks, has become a major notification mechanism, especially for point-to-point communications. Numerous "bulletin board" systems and the USENET newsgroup distribution network are popular broadcast notification
25 systems. But it is the advent of the World Wide Web, frequently referred to as "the Web", that has excited the imagination of thousands of consumers, providers, and entrepreneurs. The Web conveniently delivers text, images, and audio clips to

-9-

individual users. In fact, the Web can be used to distribute any sort of information that can be represented in a computer data file.

The popularity of the Web has encouraged the establishment of many Internet Access Providers (IAPs), who provide communication access to the Internet for individuals and organizations; and Internet Service Providers (ISPs), who provide various services via the Internet, such as e-mail delivery, Web site hosting, search engines, and "chat" areas. This communications and information infrastructure continues to grow at a prodigious rate. With so many individuals and organizations now having convenient and inexpensive communications access, the Internet offers a promising base for a new mode of commerce.

Electronic commerce addresses many of the disadvantages of traditional commerce. It is convenient and inexpensive to prepare and deliver e-mail to specific persons or groups of persons. Many computer systems are repositories for immense databases that are useful for commerce, and the global communication network provides a means for accessing that data. Personal computer systems and specialized software are now enabling consumers to view online product catalogs and other information that providers publish on the Web. A multitude of researchers and organizations are working out the details of payment mechanisms to allow secure monetary transactions across the Internet.

20

The Disadvantages of Today's Electronic Commerce

Even with the colossal potential of the Internet, there are still a number of problems to be solved to support the establishment of a viable virtual marketplace, especially regarding the collection and exchange of market information. The electronic form of commerce doesn't address all of the problems of traditional commerce, and it raises a number of new difficulties.

25

-10-

Here are some of the disadvantages, from the consumer's point of view, of the electronic form of commerce:

- 5 * Most search engines (Web sites that implement a capability for searching a database of information) are generic. They use general words as search keys, whereas consumers would benefit from information that is categorized by brands, product names, product category, store names, etc. Even though search engines are much faster than physical store visits, the search process is still tedious and prone to error.
- 10 * Using search engines for comparative shopping is very slow. Search engines return pointers to information sites, not the actual information. Search engines frequently return thousands of "hits", or items that partially match the search request. Consumers must sift through these hits, determine which ones are likely to be truly of interest, and contact the individual Web sites to collect the product information.
- 15 * Consumers find comparative shopping tedious because every Web site has its own format for information. It is difficult to automate comparative shopping because of the inconsistent and non-standardized data formats.
- * It is difficult for consumers to find independent opinions about product quality, comparative features, and how a provider treats other consumers.
- 20 * Every search starts from scratch. Even though some search engines now have the capability to narrow a search during an episode of use, each episode of searching starts anew requiring the searcher to enter all of the relevant keywords again. The preferences of the consumer are not retained between uses of the search engine.
- 25 * Consumer searching is not private. The search engine can collect data about who is searching and the keywords of their search. Many Web sites maintain "cookies" or "passports", that is, files that contain information about the

-11-

consumer who is searching. In effect, the consumer's searching and decision making is exposed to public view.

* Dealing directly with the provider exposes the consumer's identify and other data to the provider. When the consumer orders or purchases a product
5 from a provider's Web site, the consumer must reveal name, delivery address, credit card data, etc. Even if the consumer is merely inspecting the information available on a Web site, the site's owner can still collect data about the consumer from the consumer's browser software.

* Non-technical consumers may experience frustration in trying to
10 construct appropriate queries for search engines. The syntax rules for queries commonly use Boolean logic and special separator characters. Even when the use of one search engine is mastered, the consumer must learn yet another set of rules for constructing queries for another engine, since the various engines use different syntax rules for their queries.

* Once a search engine successfully delivers a promising Web address
15 (known as a "URL" for Uniform Resource Locator), the consumer may be disappointed to find that the URL is no longer valid. Thousands of Web pages are published and withdrawn daily, and the search engines are not always informed of the changes. Web pages customarily contain references to other Web pages
20 ("links"), and a link is not always updated when the target URL changes, especially if the target URL names a page that is published by another entity. These "broken URLs" refer to Web pages that no longer exist or have moved to another address, so the consumer can no longer reach the information.

* The style of presentation is still controlled by the provider. The
25 consumer has no useful mechanism to request, for example, only summary information about products. The avenue of presentation is also fixed. If the data is available on the Web, it must be accessed via the Web; the consumer typically

-12-

cannot arrange to have the data delivered, for example, via pager, facsimile, or cell phone display.

5 * Usually there is no direct benefit to a consumer (a "consideration") for viewing an electronic advertisement, although some companies have proposed paying consumers for reading advertisements via e-mail or other electronic delivery.

10 * Consumers have few means for publishing their own electronic advertisements for buying or selling. Most newsgroups do not accept advertisements. Even in the special newsgroups that do accept advertisements, it wastes bandwidth since most readers of the newsgroups won't be interested in a specific advertisement.

Electronic commerce also presents a number of difficulties for providers, especially in the areas of deploying advertising and gathering "market intelligence." Some of these disadvantages are:

15 * Many providers are reluctant to advertise on the Internet because of software agents that make recommendations based on price alone. For example, the BargainFinder service, a research project of Andersen Consulting, that gathers pricing data on audio compact disks has been blocked from many providers' Web sites. Without the opportunity to present other features that justify a higher price, higher-priced providers would lose sales or be forced to lower their prices (and erode their retail margins).

20 * Consumers may not be aware of a provider's Web site. A provider typically advertises the existence of a Web site via traditional means, driving up the cost of maintaining a Web presence.

25 * Very few tools have emerged to help providers judge the effectiveness of their Web sites. For example, providers cannot gather enough

-13-

information to calculate market share, since the statistics associated with competitor's Web sites are not publicly available.

* Providers have little control over the search engines that consumers use to locate products, beyond carefully choosing a few keywords. They can't
5 choose to emphasize different aspects of their products depending on the origin of the request.

* It is impossible to collect useful demographic data about many online consumers, because consumers frequently use pseudonyms to disguise identity.

10 * Due the lack of useful demographic data about online consumers, there is little guidance in targeting advertisements to potential customers.

* Even though it is easier to personalize e-mail, how does the provider determine the target audience? Many online users summarily reject unsolicited e-mail advertising, disparagingly called "spam".

15 * Many Web sites now have the capability to generate Web pages ("content") on the fly, but it is difficult for the Web server to obtain enough data about the requestor to personalize the content in a useful way.

* "Banner" advertisements placed on popular Web sites have not been particularly successful. Many online users don't "click-through" the banner to the
20 more extensive advertiser information, because the placement of such banners is not finely targeted. Some Web activity statistics indicate that only one and one-half to three and one-half percent of users click-through.

* Providers lose the goodwill of potential customers when they place advertisements in regular newsgroups.

25 * Providers still have few ways to gather real-time feedback on special promotions or specific products.

-14-

* Perhaps most critically, providers still have no method to derive useful market intelligence from the vast activity online. Providers still can't determine why consumers accept or reject offers, they can't calculate consumer demand, and they can't simulate demand based on actual demand data.

5

Other Electronic Commerce Systems

Much of the research relating to electronic commerce has been directed towards designing and implementing secure online money transactions. For example, U.S. Pat. No. 5,557,518 entitled "Trusted Agents for Open Electronic Commerce" issued 17 September 1996 describes a system wherein a customer and a merchant can exchange electronic merchandise and electronic money by using trusted electronic agents. Even when this crucial aspect of electronic commerce is satisfactorily resolved, there remains the challenge of establishing a mechanism for the presentation, gathering, and exchange of market information in a way that encourages and supports broad participation in the online market. Of the current mechanisms related to online market information, most tend to fall into two categories: search engines, and various kinds of software or electronic agents.

10

15

Various problems with search engines have already been mentioned. Each engine has a different syntax and operation, making them error-prone and tedious to use. The information that engines return may be out of date or just plain incorrect. Search engines don't generally categorize information in a format that is handy for consumers, since they cater to general information seekers. Searches cannot make use of personal information about users, since the search engine has access only to the specific query data entered by the user.

20

25

A software agent is a software entity that is capable of performing certain delegated electronic actions (including holding information) on behalf of a user or another agent. An IBM white paper, "The Role of Intelligent Agents in the

-15-

Information Infrastructure" (Gilbert, et. al., IBM Corporation, undated; also published on the Web at URL <http://www.networking.ibm.com/iag/iagptc2.html>; also hyperlinked from <http://www.raleigh.ibm.com/iag/iaghome.html>) describes three dimensions along which intelligent agents may be measured: agency, intelligence, and mobility. "Agency is the degree of autonomy and authority vested in the agent... *Intelligence* is the degree of reasoning and learned behavior... *Mobility* is the degree to which agents themselves travel through the network..." (IBM, italics author's). Software agents can be further classified along the dimension of mobility into three broad categories according to their location of execution and location of data reference: mobile agents, "wandering" agents, and local or static agents.

True mobile agents are software entities that can electronically move from one computer system to another. The software program of a mobile agent actually executes on the target computer system. Although some technology to support mobile agents is available (for example, Sun's Java and General Magic's Telescript), they have not been successful, partly because many computer "firewalls" block the entry of mobile agents for security reasons, and because the agents must be capable of operating on a number of specific kinds of computers.

A wandering agent is a software entity that resides within a single computer system and "visits" or communicates with other computer systems, frequently via the Internet. Wandering agents are being used successfully to map the Web, gathering the data that is used in the internal indexes of search engines. However, these agents are very slow in operation, especially when there are thousands of sites to visit, and some wandering agents may be blocked from accessing some sites (as the BargainFinder agent has been). As described in "Internet Agents: Spiders, Wanderers, Brokers, and 'Bots" (Cheong, Fah-Chun, New Riders Publishing, 1996), wandering agents are also used for various Web maintenance

-16-

tasks and for Web mirroring. Cheong lists and describes many instances of wandering agents. The following list of wandering agents was compiled on 26 December 1995 from Appendix G of Cheong. The purpose of each agent is excerpted by the inventor from short descriptions in Cheong.

5	<i>Name</i>	<i>Purpose</i>
	ASpider (Associative Spider)	searches for keywords
	Arachnophilia	collect documents
	Aretha	(none given)
10	CS-HKUST WWW Index Server	Resource Discovery, validate HTML
	ChURL	URL checking
	Checkbot	(none given)
	EIT Link Verifier Robot	verify links
	Emacs W3 Search Engine	Resource Discovery
15	Fish Search	Resource Discovery
	GetURL	validate links, mirroring
	HTML Analyzer	check validity of Web servers
	HTMLgobble	mirroring
	Harvest	Resource
20	InfoSeek Robot	collect information for database
	JumpStation Robot	Resource Discovery
	Katipo	look for changed documents
	Lycos	information retrieval and discovery
	MOMspider	maintenance of distributed hypertext
25	Mac WWWorm	keyword searching
	NHSE Web Forager	Resource Discovery
	NIKOS	Resource Discovery

-17-

	NorthStar Robot	textual analysis, indexing
	Open Text Corporation Robot	(none given)
	Peregrinator	indexing
	Python Robot	(none given)
5	RBSE Spider	Resource Discovery
	SG-Scout	Resource Discovery
	Scouter	Resource Discovery
	Spry Wizard Robot	Resource Discovery
	TITAL	Resource Discovery
10	Tarspider	mirroring
	Tel W3 Robot	validate links
	TkWWW Robot	find logically related pages
	W4 (World Wide Web Wanderer)	measure growth in Web
	WM32 Robot	Resource Discovery, validate links
15	WWW - World Wide Web Worm	indexing
	WebCopy	mirroring
	WebCrawler	Resource Discovery
	WebLinker	traverses Web converting URN->URL
	WebWatch	validate HTML
20	Webfoot Robot	(none given)
	Weblayers	validate, cache, maintain links
	Websnarf	mirroring
	Webwalk	Resource Discovery, validate links, mirroring

25 A local or static agent is a software entity that operates within a single logical computer system, accessing data local to that system. Clearly this kind of agent is of limited usefulness for the electronic marketplace, since, a local agent would not have access to the variety of data that is necessary for a thriving

-18-

marketplace. However, if there were a mechanism to incorporate data from many sources and to provide access to a broad base of users, local agents could be extremely useful.

5 These technologies have not been used to collect market information that providers can use to quantify consumer demand or to help gain customers at reduced cost. Consumers are hesitant to use some of these technologies because of privacy concerns. There isn't a practical mechanism for the user to instruct an agent to "keep looking" if the immediate search fails or is only partially
10 successful. Although these technologies may be useful for the electronic marketplace, additional mechanisms are required for practical, ubiquitous electronic commerce.

The Fundamental Problems to be Solved to Enable Electronic Commerce

15 An electronic marketplace, just like a traditional marketplace, must support the basic process of commerce: offers to sell or buy are made, offers are accepted, and considerations (payments) are paid. If the basic process does not work, there is no marketplace. A viable marketplace must also address side-effects of commerce such as issues of security, privacy, and confidence or trust; otherwise, even if the basic process works, consumers and providers will not feel
20 comfortable enough to participate in the marketplace.

A practical and viable electronic marketplace involves the exchange of market information, as well as the more obvious trading for goods and services. From a consumer's point of view, shopping is a means of gathering data about goods and services offered. This data is used by the consumer to compare and rank
25 offerings and to make decisions about purchases. From a provider's point of view, consumer shopping is an opportunity to gather data about consumer needs and

-19-

interests. This data is used by the provider to improve product and service offerings.

For consumer, the fundamental problems relating to the flow of market information in electronic commerce, still to be solved, are:

- 5 * Consumers need help in gathering information about available goods and services, quickly and with a minimum of fuss;
- * Consumer privacy needs to be assured while gathering information about available goods and services; and
- * Consumers need a mechanism for ongoing, autonomous searches
10 for information about available goods and services, searches that continue even when the consumer is not "on-line".

For providers, the fundamental problems relating to the flow of market information in electronic commerce, still to be solved, are:

- 15 * Providers need to target advertising information to truly interested consumers without disturbing the privacy of those consumers;
- * Providers need to be able to quantify consumer demand, both offline and in real-time, using historical and current data;
- * Providers need to be able to determine reasons for sales and lost sales; and
- 20 * Providers need a source of more accurate market data to serve as input to present and yet-to-be-developed market analysis methods.
- * Solutions to these problems for providers must be cost-effective.

Objects and Advantages

- 25 * The principal object of the present invention is to provide a system that facilitates the gathering and exchange of market information in support of

-20-

electronic commerce. The attainment of this object has many advantages for both consumers and providers.

Several objects and advantages of the invention for consumers are the following.

- 5 * A consumer's identifying and private information is protected from disclosure as the consumer gathers information about available products.
- * Consumers can identify suitable products more easily, with less expenditure of effort, because the product information is presented in a consistent manner.
- 10 * Consumers can identify suitable products more quickly.
- * Consumers can use advocate recommendations and evaluations in deciding between competing products.
- * Consumers have more control over the presentation of advertising information. They can control what information is permitted to be delivered, when
15 the information is delivered, and what devices are used for delivery.
- * Consumers have a standardized mechanism for receiving considerations from advertisers in exchange for allowing delivery of advertisements and other provider information.
- * Consumers can launch ongoing searches for products, and the
20 searches can continue even when the consumer is not online.
- * Consumers use search engines that have data that is more up-to-date.
- * Consumers access search engines that are easier to use, especially
 for non-technical users.
- 25 * Consumers can place their own want-to-buy and want-to-sell advertisements.

-21-

Several objects and advantages of the invention for providers are the following.

- * A provider's privacy is protected while searching for potential customers or surveying competitor's offerings.
- 5 * Providers can use demographic and preference data that is more up-to-date.
- * Providers own and control the information about the products they offer, even when the information resides within the system contemplated by the invention.
- 10 * Providers can directly contact more consumers that are ready to buy.
- * Providers can target consumers more economically.
- * Advertising may have higher success rates since the targeted consumers have expressed an interest in the product.
- 15 * Providers can personalize special offers based on previous buying habits and future intent of the consumer.
- * Providers have a mechanism for quantifying consumer demand.
- * The mechanism for quantifying consumer demand is based on actual consumer buying data, both historical and current.
- 20 * The mechanism for quantifying consumer demand uses data based on individual buying decisions, not merely aggregate or estimated data.
- * Providers can quantify demand in real-time.
- * Providers have a mechanism for discovering the reasons for lost sales.
- 25 * Providers can provide a consideration to consumers for viewing advertisements and other notices.

-22-

* Providers can receive feedback in real-time about the success of promotions.

* Providers have access to market data based on individual consumer needs that can be used to simulate demand in various scenarios.

5 Several objects and advantages of the invention for all participants in the system are the following.

* A robust and thriving electronic marketplace may lower the amount of wasted paper and energy for the delivery of printed matter.

* The system provides results faster than mobile or wandering agents.

10 * Many industries and providers can participate in the system.

* Information used by both consumers and providers is more up-to-date.

Referring to the fundamental problems of the flow of market information in electronic commerce, the fundamental objects of the system for consumers are:

15 * to assist consumers in gathering market information quickly and easily;

* to protect consumer identity and private information while gathering market information; and

* to assist consumers in performing ongoing searches.

20 Referring to the fundamental problems of the flow of market information in electronic commerce, the fundamental objects of the system for providers are:

* to assist providers in targeting information delivery to interested consumers;

25 * to assist providers in quantifying consumer demand, both offline and in real-time, using historical and current data;

* to assist providers in determining reasons for sales and lost sales;

and

-23-

- to assist providers by furnishing a huge base of accurate market data based on actual consumer activity to serve as input to present and future market analysis methods; and

- to assist providers in gathering and analyzing market information in a cost effective way.

Further objects and advantages of this invention will become apparent from a consideration of the drawing and ensuing description.

Brief Description of the Drawing

10 The invention is described in greater detail below with reference to the attached drawing. In the drawing, closely related figures have the same number but different alphabetic suffixes.

FIGURE 1 is a schematic diagram showing the basic interaction of several kinds of agents within an agent system.

15 FIGURE 2 is a schematic diagram of an agent system and its major components.

FIGURE 3A is a topological diagram showing an example arrangement of processors in an agent system.

FIGURE 3B illustrates the functional components of a processor.

20 FIGURES 4A - 4D illustrate the functional components of Personal Agents.

FIGURE 5A illustrates the data components of a Preference Datum.

FIGURE 5B shows example Preference Data.

FIGURE 6 illustrates the functional components of Decision Agents.

25 FIGURE 7 illustrates the functional components of Demand Agents.

FIGURES 8A - 8C illustrate the functional components of Markets.

-24-

FIGURES 9A - 9C illustrate the functional and data components of a Product Database.

FIGURE 9D shows example Product Template Entries with example values.

5 FIGURE 10 illustrates the data components of an Ad.

FIGURE 11 is a flow diagram representation of an overall method for searching for products.

FIGURES 12A - 12B comprise a flow diagram representation of a method for composing a Decision (product search) query.

10 FIGURE 13 is a flow diagram representation of a method for creating a Decision Agent.

FIGURE 14 is a flow diagram representation of a method for accepting a new Decision Agent into a Market.

15 FIGURE 15 is a flow diagram representation of a method for performing a Decision search.

FIGURE 16 is a flow diagram representation of a method for performing an Immediate search portion of a Decision search.

FIGURE 17 is a flow diagram representation of a method for completing a Decision search.

20 FIGURE 18 is a flow diagram representation of a method for performing an Extended search portion of a Decision search.

FIGURE 19 is a flow diagram representation of a method for delivering Decision search results to the consumer.

25 FIGURE 20 is a flow diagram representation of a method for expiring a Decision Agent that has completed its task.

FIGURE 21 is a flow diagram representation of an overall method for quantifying demand.

-25-

FIGURES 22A - 22B comprise a flow diagram representation of a method for composing a Demand query.

FIGURE 23 is a flow diagram representation of a method for creating a Demand Agent.

5 FIGURE 24 is a flow diagram representation of a method for accepting a new Demand Agent into a Market.

FIGURE 25 is a flow diagram representation of a method for performing a Demand search.

10 FIGURE 26 is a flow diagram representation of a method for performing a search for current demand.

FIGURE 27 is a flow diagram representation of a method for performing a search for historical demand.

FIGURE 28 is a flow diagram representation of a method for delivering Demand search results to a provider.

15 FIGURE 29 is a flow diagram representation of a method for expiring a Demand Agent that has completed its task.

FIGURE 30 is a flow diagram representation of an overall method for placing an Ad in a Market.

20 FIGURES 31A - 31B comprise a flow diagram representation of a method for composing an Ad.

FIGURE 32 is a flow diagram representation of a method for creating an Ad.

FIGURE 33 is a flow diagram representation of a method for accepting a new Ad into a Market.

25 FIGURE 34 is a flow diagram representation of a method for expiring an Ad that has reached its expiry time.

-26-

FIGURE 35 is a flow diagram representation of an overall method for targeting a group of consumers.

FIGURE 36 is a flow diagram representation of a method for selecting Personal Agents that represent targeted consumers.

5 FIGURE 37 is a flow diagram representation of an overall method for rejecting unsolicited messages.

FIGURE 38A-B is a flow diagram representation of an overall method for simulating demand and for replaying demand.

10 FIGURE 39 is a schematic representation of a Web page used to "login" to a agent system.

FIGURE 40 is a schematic representation of an example Web page used to specify search criteria when composing a Decision query in an example consumer electronics Market.

15 FIGURE 41 is a schematic representation of an example Web page used to specify search criteria when composing a Decision query in an example automobile Market.

FIGURE 42 is a schematic representation of an example Web page used when composing an advertisement for a television set.

20 Summary of the Invention

The present invention contemplates a system for enabling the collection of market information, especially data needed to quantify various kinds of consumer demand, while protecting the particular identity and privacy of consumers. Consumers, because their identity is protected, feel secure in using the system, thereby generating market data as a by-product of their shopping activities. Providers can query and analyze this market data in many ways, including the calculation of actual instantaneous and historical consumer demand for products

25

-27-

and product categories. This kind of market data is not available by any other means.

The system comprises a "virtual marketplace" in which various kinds of agents represent human owners. Although the system has access to data outside its boundaries, the system behaves as a logically single system. The agents of the system need not travel outside the system to perform their tasks.

Consumers and providers both may place sell and buy advertisements (ads) in the marketplace. Providers can target groups of consumers to receive special messages such as special offer ads.

Consumers and providers are each represented in the system by Personal Agents. A Personal Agent stores and learns the preferences of its human owner and arranges for delivery of messages to the owner according to the owner's desired delivery times and desired delivery devices. The Personal Agent ensures that private or identifying data about the owner is never revealed without authorization to other agents in the system.

Consumers use Decision Agents to gather the information that helps consumers make purchasing and usage decisions. Decision Agents can search for ads meeting various criteria, and order the matching ads according to the consumer's preferences.

Providers use Demand Agents to assist with market analysis of various kinds of demand and to target consumers. Demand Agents can target consumers based on consumer preferences, demographics, and shopping activity.

Description of the Preferred Embodiment

Within this description, the term "product" is understood to include all kinds of merchandise, tangible and intangible goods, services, intellectual

-28-

property, information, electronic merchandise, etc. whether offered for a price, offered for barter, or offered for free.

The term "consumer" is understood to mean a user of the system who is acting to find information about or purchase products offered by other users of the system. A consumer is typically an individual. The term "provider" is understood to mean a user of the system who is acting to present information about products and/or sales offers to other users of the system. The term "provider" includes manufacturers, retailers, wholesalers, distributors, etc. When the term "consumer" is used in this description, it is understood to mean a user of the system acting in the role of a consumer. When the term "provider" is used, it similarly is understood to mean a user of the system acting in the role of a provider. The term "user" is understood to apply in a context where the particular role is unimportant.

Agent Interaction

15

Referring now to the drawing, and in particular to Figure 1, an Agent System 10 is shown schematically, together with some of its internal components, in order to illustrate the basic interaction between several system components in a preferred embodiment of the system. Two users, a Consumer 20 and a Provider 21, who are not part of the invention, are shown to illustrate their relationship to Agent System 10. Many components of the system are not illustrated in this figure, in order to focus attention on the basic interaction, which enables the generation and the retrieval of market data.

Referring to the left side of the figure, actions of Consumer 20 generate market data. Consumer 20 controls a Consumer Personal Agent 12 that represents the Consumer to the system. The Consumer Personal Agent is capable of creating a Decision Agent 14 to carry out a search, within a Market 18, for products that

25

-29-

satisfy certain constraints and preferences. For example, a Consumer might query for the local retailers that carry a certain brand of sports shoes. Decision Agent 14 gathers data without knowing, and therefore without revealing, the identity of the Consumer 20. Both Decision Agent 14 and Market 18 store data about the search.

5 Decision Agent 14 returns a set of product recommendations, which Consumer Personal Agent 12 further filters and orders according to Consumer preferences before presenting to Consumer 20.

Referring to the right side of the figure, actions of Provider 21 retrieve and analyze market data. Provider 21 (merchant, service provider, etc.) controls a

10 Provider Personal Agent 13 that represents the Provider to the system. The Provider Personal Agent is capable of creating a Demand Agent 16 to collect data, from a Market 18, about consumer demand. For example, the Provider might query for the number of consumers that are currently searching, or have searched within the past 24 hours, for a certain brand of sports shoes. Demand Agent 16

15 accesses data stored in the Market 18 and in active and expired Decision Agents 14. Demand Agent 16 returns a response for the query to the Provider Personal Agent 13, which uses the Provider's preferences to determine how to present the retrieved data to the Provider 21.

Agent System 10 contains different Markets 18 for various categories of

20 products and services. The various kinds of agents and the markets are software components. These components are more fully described in conjunction with other figures. In a preferred embodiment, the software components utilize but are not limited to conventional object-oriented technology, distributed object-oriented technology, object-oriented database technology, relational database technology,

25 general Internet communication technology, World Wide Web (WWW or Web) technology, and electronic mail (e-mail) technology.

-30-

Agent System Schematic

Referring to Figure 2, there is shown a more detailed schematic diagram of a preferred embodiment of the invention. Agent System 10 is a combination of hardware and software components; the components, taken together, are considered to be a single, logical system, regardless of the physical topology of the hardware components. It is expected that the physical topology of Agent System 10 will change over time, especially for the addition of processing units to distribute the load as more users participate in the system.

Two users of Agent System 10 (not part of the invention) are shown schematically, a User 21 (a Provider) and another User 20 (a Consumer), to indicate their relationship to the Agent System. When a User 20 or 21 uses Agent System 10 to present information about products and/or sales offers to other users of the system, that User is referred to as a Provider. For example, merchants, distributors, retailers, wholesalers, etc. fall in this category. When a User 20 or 21 uses Agent System 10 to find information about or purchase products offered by other users of the system, that User is referred to as a Consumer. In this document, the term "products" comprises products, services, tangible goods, intellectual property, etc. Persons and organizations are registered to become users of Agent System 10, and each user is authorized to perform certain functions. Not all functions of Agent System 10 are permitted to every user.

Each User of Agent System 10 has a means of communicating with the system, as indicated by Communication Device 22 or 23 of each User 20 or 21 respectively, and may have multiple means and devices. These communication devices can be any device capable of communicating over the Internet (such as personal computers with Web browser and/or e-mail software), other devices capable of operating with computer control (such as facsimile machines and

-31-

5 pagers), and other means of transferring data and commands between the User and the Agent System. A User may also store data on various devices outside Agent System 10, as indicated by Provider Remote Data 25 and Consumer Remote Data 24, provided that such data can be accessed by the system via (at least one of) the User's Communication Device(s). In this context, "remote" means located outside the direct control of Agent System 10.

10 Agent System 10 contains software agents representing both consumers and providers. Each User controls a Personal Agent 12 or 13 (PA) that coordinates the communication of the User with the other parts of Agent System 10. A User establishes a communication session with User's PA using conventional authentication means appropriate to the Communication Device 22 or 23.

15 A Personal Agent Search Engine 26 maintains indexes over preference data and demographic data of all Personal Agents, so that users may query to identify a set of Personal Agents whose users have certain characteristics or preferences. However, private data about the user (name, address, etc.) is not maintained in Personal Agent Search Engine 26.

20 Continuing to refer to Figure 2, an Agent Marketplace 28 within Agent System 10 provides a means for various agents to interact on behalf of their owners. Consumer's Decision Agents 14 and Provider's Demand Agents 16 are "launched" into the marketplace to perform their delegated tasks. Agent Marketplace 28 comprises a number of Markets 18, which may be either General Markets 18a, in which all Users may launch specialized agents or place advertisements, or Restricted Markets 18b, in which only authorized Users may launch specialized agents (such as Decision Agents 14 or Demand Agents 16) or place advertisements. General Markets 18a correspond to conventional broad product categories; some examples are: Home Appliances, Office Supplies, Groceries, Consumer Electronics, Residential Real Estate, Commercial Real

-32-

Estate, etc. Restricted Markets are used, for example, by wholesalers to restrict access to special prices to distributors. The basic structure of both General Markets 18a and Restricted Markets 18b is the same, and is described in conjunction with Figure 8A below. Various specialized agents are described in conjunction with other Figures. Agents and other components operating in Agent Marketplace 28 have access to a Product Database (Product DB or PDB) 32.

Agent Marketplace 28 contains a number of Market Navigation Aids 34 to help users find the appropriate Market 18 in which to search or place ads. These aids might include keyword searches, word alias searches, hierarchical browsers of market layouts, etc..

Agents and other components of Agent System 10 record and access system history data (records of searches, transactions, etc.) in System History Data 36 component. Most of the system history is more conveniently accessed through logs and archives located within various functional components, but System History Data 36 maintains the "master" copy.

System Administrator users (not shown) perform or supervise various conventional maintenance functions for Agent System 10, such as performing backups, adding new product data, redistributing functions between processors for load balancing, etc.

Agent System Topology

Referring to Figure 3A, there are shown hardware components of Agent System 10 from a topological point of view. Agent System 10 may have any convenient hardware topology; Figure 3A is intended as an example. The hardware of Agent System 10 comprises a number of Processors 38, each capable of communicating with other Processors. As illustrated, normally a group of

-33-

Processors is clustered together, with one or more Processors in a group configured to communicate with one or more Processors in other clusters. The clusters may be geographically dispersed, and Processors within a cluster may be geographically dispersed. The actual number of Processors and their topology will
5 change over time, to support additional capacity, load balancing, and ease of administration.

The various functional components of Agent System 10, described in conjunction with several Figures, reside on one or more Processors 38, and may be duplicated to reside on one or more Processors 38 simultaneously. The
10 distribution of the functional components across the various Processors 38 will change over time, to support additional capacity, load balancing, and ease of administration.

Referring to Figure 3B, a Processor 38 comprises the functional components of:

15 an Object Server function 40,
a Communications function 42,
a Human/Machine Interface function 44,
a Messaging function 46,
and a Persistence function 48.

20 These functional components of Processor 38 are available for use by any software component of Agent System 10 that resides on Processor 38.

A Object Server function 40 executes the software objects that comprise the various software functional components of Agent System 10, for example, the various agents, the markets, the data repositories, and lower level utility software
25 objects (not shown). A particular Object Server 40 need not execute every kind of object; for example, some Object Servers 40 may contain only Personal Agents 12

-34-

or 13 and their related data repositories, but will not contain any Market 18 objects.

5 A Communications function 42 supports communications between Processor 38 and the outside world. Such communications may be wired or wireless, broad or narrow band, so long as the Processors 38 use compatible communications. Communications function 42 sets up the connection between two Processors 38, or connects a Processor 38 to a network for indirect connection to another Processor 38 or to a User's Communication Device 22 or 23.

10 A Human/Machine Interface function 44 provides the look and feel of Processor 38. It could include a keyboard, mouse, pen, voice, touch screen, icons, menus, etc. Human/Machine Interface function 44 communicates with other functions in Processor 38. In some situations, a Human/Machine Interface function may not be necessary, for example, when a Processor 38 communicates only with other Processors 38 but not with a User's Communication Device 22 or
15 23.

A Messaging function 46 routes messages between software objects executing on various Processors 38.

20 A Persistence function 48 manages storage of data belonging to the various software objects that reside on the Processor 38. The actual data is stored on conventional storage devices (not shown), such as computer disks.

Personal Agent

25 Personal Agent 12 or 13 is the point of contact between a user and the Agent System 10. Personal Agent 12 or 13 acts as an electronic "butler" or assistant, accepting requests from the user, delegating tasks to other agents in the system, and arranging for responses from various agents to the user to be delivered

-35-

at a time and in a manner that is convenient for the user. Consumer Personal Agent 12, via its internal functions, maintains the user's preferences and other data about the user, some of which is protected from unauthorized access.

Referring to Figure 4A, a Personal Agent 12 or 13 comprises the

5 functional components of:

- a Unique identification (ID) 50,
- an Owner Manager 52,
- a Preference Manager 54,
- a Delivery Manager 56,
- 10 an Individual Firewall 58,
- a Decision Agent Manager 60,
- a Demand Agent Manager 62,
- an Ad Manager 64,
- a Target Manager 66, and
- 15 a Consideration Account 67.

A Unique ID function 50 maintains an identifier that uniquely identifies this agent within Agent System 10. Unique ID 50 is generated automatically when the agent is created, and is never reused to identify a different agent, even if the original agent ceases to exist within the Agent System. Unique ID 50 carries no

20 information that reveals the human "owner" of this agent. Unique ID 50 is used to address messages to the agent.

An Owner Manager function 52 maintains data about the human "owner" of the agent, i.e. the user that controls this Personal Agent 12. This data includes the user's name, postal addresses, e-mail addresses, telephone and fax numbers,

25 etc. This data is always protected by an Individual Firewall 58; it is never revealed to other agents, and is used only by components of Personal Agent 12 to deliver messages, for system invoicing, etc.

-36-

A Preference Manager function 54 maintains data about the preferences of the user. Preferences indicate items of interest to the user, such as favorite brands, interest in sports, etc. Within Agent System 10, preference data also includes "demographic" data. Demographic data indicates facts about the user, such as whether the user is a homeowner, the user's gender, the user's age group, etc. Although marketing industry usage of the term "demographics" may include a person's name, address, or other identifying data, a Preference Manager's demographic data does not include data that identifies the particular user. Preference data may be entered manually by the user using, for example, a form on a Web page, or data may be loaded by a System Administrator. Preferences may also be updated automatically by the system as, for example, when the user instructs the system to "remember" a product brand name from a product search. Preference Manager 54 uses preference data to order search results, so that items that are more likely to be preferred by the user will be displayed first when the results are delivered to the user. Referring now to Figure 5A, each preference datum 68 comprises not only a value 72, but also a key 70 for ease of searching. Referring to Figure 5B, a small sample of preference data illustrates the kind of data that might be used. A particular user typically will have much more preference data. Some values are shown as "rank m in n" to illustrate that ranking data may also be stored. The specific keys of any particular set of preference data depends on what the user has entered, etc. Only keys that are relevant to a particular user are included in that user's preferences, and the specific data maintained will change over time.

Referring again to Figure 4A, a Delivery Manager function 56 accepts all messages, generated by agents or other components of the system, that are directed to the user, and delivers those messages according to the user's desired delivery time and delivery media. Default delivery time and delivery media are specified as

-37-

part of the user's preferences (maintained by Preference Manager 54). Individual messages may also have a specified delivery time and delivery media that overrides the defaults. Delivery Manager 56 establishes communication with the user's Communication Device 22 or 23 to effect delivery. Messages may be sent
5 to multiple devices if the user so desires. Delivery Manager 56 queues messages that are to be delivered at a future time.

Delivery Manager 56 also rejects unsolicited messages, unless the message meets the user's preferred criteria as maintained by Preference Manager 54. Delivery Manager 56 sends a rejection message in reply to the original sender of a
10 rejected message. The rejection message indicates why the original message was rejected, so that the sender may gather quantifiable feedback.

An Individual Firewall function 58 mediates all access to the data that is maintained by the various internal functions, ensuring that only authenticated and authorized agents and users can access private data.

15 A Decision Agent Manager 60 assists the user with the creation and management of Decision Agents 14. Referring now to Figure 4B, a Decision Agent Manager 60 comprises the functional components:

- 20 a Decision Composer 74,
- a Decision Agent Factory 76,
- a Decision Agent Tracker 78,
- and a Decision Agent Archive 80.

Continuing to refer to Figure 4B, a Decision Composer 74 assists the user in composing queries to be executed by Decision Agents. Decision Composer 74 retrieves a Product Template 174 (described later in conjunction with Figure 9B)
25 for a particular product from a Market 18 in which the user wishes to search, present instructions to the user for completing Product Template 174 to describe the object of the search, and produces the appropriate query.

-38-

A Decision Agent Factory 76 creates a new Decision Agent 14 to carry out a query.

A Decision Agent Tracker 78 enables the user to monitor the activities of Decision Agents 14 that are active, i.e., agents that have not completed their tasks. Decision Agent Tracker 78 also enables the user to cancel an active Decision Agent 14 before its scheduled expiry time.

A Decision Agent Archive 80 stores and accesses Decision Agents 14 that are expired, i.e., agents that have completed their tasks, whether successfully or not. For example, if a Demand Agent 16 needs more detailed data about a query than is stored in a Query Logger 136 of a Market 18, it can access the details of the related Decision Agent 14 through Decision Agent Archive 80.

Referring back to Figure 4A, a Demand Agent Manager function 62 assists the user with the creation and management of Demand Agents 16. Only users of Agent System 10 who are authorized to launch Demand Agents will have a Demand Agent Manager 62 as part of their Personal Agent. Referring now to Figure 4C, a Demand Agent Manager 62 comprises the functional components:

- a Demand Composer 82,
- a Demand Agent Factory 84,
- a Demand Agent Tracker 86,
- and a Demand Agent Archive 88.

These components provide functions similar to the analogously-named functional components of Decision Agent Manager 60 (described above), except that Demand Agent Manager 62 components work with Demand Agents 16 instead of Decision Agents 14.

Referring back to Figure 4A, an Ad Manager function 64 assists the user in placing advertisements. Referring now to Figure 4D, an Ad Manager 64 comprises the functional components of:

-39-

an Ad Composer function 90,
an Ad Delivery function 92,
an Ad Tracker function 94,
and an Ad Archive function 96.

5 An Ad Composer function 90 retrieves a Product Template 174 (described later in conjunction with Figure 9B) for a particular product from a Market 18 in which the user wishes to advertise, presents instructions to the user for completing Product Template 174 to describe the product, and produces a new Ad 186 (see Figure 10).

10 An Ad Delivery function 92 delivers Ad 186 to the desired destination. For consumer users, Ad 186 is delivered to Market 18, where it is accessible to other agents in the system. For provider users, Ad 186 may be delivered to Market 18 (just like a consumer ad); or Ad 186 may be delivered (as a special offer) to a set, determined by a Demand Agent 16, of consumer Personal Agents 12.

15 An Ad Tracker function 94 monitors the activity of Ad 186, including any messages received in response to the Ad, until Ad 186 expires or is canceled by the user. Ad Tracker function 94 enables the user to cancel an Ad 186 before its scheduled expiry time.

An Ad Archive function 96 stores and access Ads 186 that are expired.

20 Referring again to Figure 4A, a Target Manager function 66 assists the user in identifying Personal Agents to which targeted ads may be delivered. Target Manager 66 can identify Personal Agents based on preferences, demographic characteristics, and Decision Agent activity. Target Manager 66 does not have access to private data of consumer Personal Agents 12 such as name, address, etc.

25 A Consideration Account function 67 maintains a "consideration" account for the user. When the user earns a consideration by, for example, viewing a directly delivered advertisement or message, or completing a marketing survey,

-40-

the consideration amount is credited to Consideration Account 67. The account is denominated in a convertible exchange media such as electronic cash tokens.

Decision Agent

5

Referring to Figure 6, a Decision Agent 14 comprises the functional components of:

- a Unique ID 98,
- a Personal Agent Reference 100,
- 10 a Market Reference 102,
- an Expiry function 104,
- a Query 106,
- a Response Manager 108,
- and a Log function 110.

15 A Decision Agent 14 acts on behalf of a consumer user, as instructed by the consumer's Personal Agent 12, to search out and collect information from Agent System 10 that helps the consumer make purchasing and usage decisions. A consumer may have multiple Decision Agents 14 active within the Agent System 10 at any time. For example, a consumer may have one Decision Agent 14
20 searching for a good buy on a certain sports shoe, and have another Decision Agent 14 searching for a television set with special features.

A Unique ID function 98 maintains an identifier that uniquely identifies this agent within Agent System 10. Unique ID 98 is generated automatically when the agent is created, and is never reused to identify a different agent, even if the
25 original agent ceases to exist within Agent System 10. Unique ID 98 carries no information that reveals the human "owner" of this agent. Unique ID 98 is used to address messages to the agent.

-41-

A Personal Agent Reference 100 holds a copy of the Unique ID 50 of consumer Personal Agent 12 that controls this Decision Agent 14. Personal Agent Reference 100 is used to address messages to the controlling Personal Agent 12.

5 A Market Reference 102 indicates in which Market 18 that Decision Agent 14 should search.

An Expiry function 104 indicates how long Decision Agent 14 should continue searching. Expiry 104 may indicate either that the search should be performed and the responses returned immediately (an "immediate search"), or that the search should continue for a specific period of time, for example, one
10 week, with responses being returned periodically during that time (an "extended search").

A Query 106 describes the product or product category for which to search. Query 106 includes data from Product Template 174 completed by the consumer and relevant data from the consumer's preferences, as assembled by
15 Decision Agent Factory 76 of the consumer's Personal Agent 12.

A Response Manager 108 receives search results and returns them to the consumer's Personal Agent 12.

A Log function 110 stores records of the activities of Decision Agent 14. These records may be consulted later, for example, by a Demand Agent 16 that is
20 calculating historical demand for a product.

Demand Agent

Referring to Figure 7, a Demand Agent 16 comprises the functional
25 components of:

- a Unique ID 112,
- a Personal Agent Reference 114,

-42-

a Market Reference 116,
a Datetime Range function 118,
a Demand Query 120,
and a Log function 122.

5 A Demand Agent 16 acts on behalf of a provider user, as instructed by the
provider's Personal Agent 13, to search out and collect information from the
Agent System 10 that helps the provider quantify consumer demand and helps
target specialized advertisements to a group of consumers. A provider may have
multiple Demand Agents 16 active within Agent System 10 at any time. For
10 example, a provider may have one Demand Agent 16 calculating historical
demand over the past month for a certain model of sports shoe, and have another
Demand Agent 16 searching for consumers who have purchased sports shoes in the
past month to receive ads for sports socks.

 A Unique ID function 112 maintains an identifier that uniquely identifies
15 this agent within Agent System 10. Unique ID 112 is generated automatically
when the agent is created, and is never reused to identify a different agent, even if
the original agent ceases to exist within Agent System 10. Unique ID 112 carries
no information that reveals the human "owner" of this agent. Unique ID 112 is
used to address messages to the agent.

20 A Personal Agent Reference 114 holds a copy of the Unique ID 50 of
provider Personal Agent 13 that controls this Demand Agent 16. Personal Agent
Reference 114 is used to address messages to the controlling Personal Agent 13.

 A Market Reference 116 indicates in which Market 18 or Markets 18 that
Demand Agent 16 should search.

25 A Datetime Range function 118 indicates that demand should be quantified
over the date/time range specified; i.e., only Decision Agents 14 that were (or are)

-43-

active during the datetime range specified should be searched when quantifying demand.

A Demand Query 120 describes a product or product category query that can be matched against the Queries 106 of Decision Agents. Demand Query 120 is similar to a Decision Agent's Query 106, except that a Demand Query 120 is matched against other queries (Decision Agent Queries 106), whereas a Decision Agent's Query 106 is matched against product advertisements. Demand Query 120 causes the selection of Decision Agents 14 whose queries are searching for certain products or product categories.

A Log function 122 stores records of the activities of Demand Agent 16 for later consultation by other components of Agent System 10.

Market

Referring back briefly to Figure 2, recall that there are a variable number of Markets 18 within Agent System 10. The Markets 18 are of two basic kinds, General Markets and Restricted Markets, which have similar structure. Referring now to Figure 8A, a Market 18, of either the General or Restricted kind, is comprised of various functional components:

- a Product Listing function 124,
- a Cross Reference (Xref) Manager function 126,
- a Sell Ad Manager function 128,
- a Buy Ad Manager function 130,
- an Active Demand Agent Manager function 132,
- a Template Dispenser function 134,
- a Query Logger function 136,
- an Historical Demand Search Engine function 138,

-44-

some number of Remote Database Adaptors 140,
and an Authorization function 142.

Market 18 corresponds to a traditional broad product category, such as Home Appliances, Office Supplies, Groceries, Consumer Electronics, Residential
5 Real Estate, Commercial Real Estate, etc. The major purposes of a Market 18 are to maintain the advertisements (ads) for products of that market, to provide the capability for specialized agents to search the advertisements, and to collect data about searches for later demand calculations.

10 A Product Listing function 124 maintains a list of the products that can be advertised in this market. Each product references detailed product data that is kept in a Product Database (PDB) 32 described in conjunction with Figure 9A.

Referring again to Figure 8A, a Cross Reference (Xref) Manager function 126 maintains, upon command from a System Administrator, cross references to other Markets 18 that carry similar products.

15 A Sell Ad Manager function 128 accepts advertisements of offers to sell that are submitted by users' Personal Agents 12. A Buy Ad Manager function 130 accepts advertisements of offers to buy that are submitted by users' Personal Agents 12. Both consumers and providers may place ads for selling or buying in a Market 18.

20 The structure of both Sell Ad Manager 128 and Buy Ad Manager 130 are similar, the difference being the kind of advertisements that are accepted. Referring to Figure 8B, either kind of Ad Manager comprises the functional components:

25 an Ad Indexing function 144,
an Active Ads function 146,
a Future Ads function 148,
an Expired Ads function 150,

-45-

an Active Decision Agent Manager function 152,

and an Expired Decision Agent Manager function 154.

Continuing to refer to Figure 8B, an Ad Indexing function 144 maintains indexes for quick searching of the ads by, for example, brand name, UPC code,
5 product name, vendor name, etc.

An Active Ads function 146 maintains the ads that are currently active. As each new add is accepted by Active Ads function 146, an Active Decision Agent Manager 152 (see below) is notified so that pending searches can be matched against the new advertisement.

10 A Future Ads function 148 maintains ads that have been submitted to the Market, but are not yet active because their effective datetime has not yet been reached. These ads are moved to Active Ads when their starting datetime is reached.

An Expired Ads function 150 maintains an archive of ads that have
15 expired, for later analysis by other components.

An Active Decision Agent Manager function 152 maintains a list of all Decision Agents 14 that are currently searching this Market 18 for products. Provider's Demand Agents 16 refer to Active Decision Agent Manager 152 while calculating current (or instantaneous) demand.

20 Referring to Figure 8C, Active Decision Agent Manager 152 comprises a number of functional components:

an Immediate Agents function 156,

a Basic Search Engine function 158,

a Pending Agents function 160,

25 an Incremental Search Engine 162,

and a Current Demand Search Engine 164.

-46-

An Immediate Agents function 156 keeps track of Decision Agents 14 that are performing an immediate search. An immediate search is a search that is to be performed and results returned as soon as practical. Immediate Agents function 156 uses a Basic Search Engine function 158 to execute the query associated with each incoming Decision Agent 14. For example, a Decision Agent 14 might present a query to find a certain brand of sports shoe within a certain price range. Basic Search Engine 158 consults Active Ads 146 to match the query against all active ads. After the immediate search is complete, if there will not be an extended search, Active Decision Agent Manager 152 moves the Decision Agent 14 to Expired Decision Agent Manager 154 (see below). If there will be an extended search, Active Decision Agent Manager 152 delivers the Decision Agent 14 to Pending Agents 160.

Continuing to refer to Figure 8C, a Pending Agents function 160 keeps track of Decision Agents 14 that are performing an extended search. An extended search is a search that remains active for an extended but specific period of time. Results from an extended search may be returned periodically during the time that the search remains active. Pending Agents 160 is notified by Active Ads 146 when a new advertisement enters the market, and uses an Incremental Search Engine 162 to match each new advertisement against the queries of the pending extended Decision Agents 14. In this way pending Decision Agents 14 are matched against ads, especially limited time special offer ads of providers, that enter the system later than the Decision Agent. Pending Agents 160 also arranges to expire Decision Agents 14 at the end of their expiry time, and to move them to Expired Decision Agent Manager 154 (see below).

A Current Demand Search Engine 164 matches demand queries of Demand Agents 16 against the queries of Decision Agents 14 that are residing in Immediate

-47-

Agents 156 or Pending Agents 160, to identify active Decision Agents 14 that are searching for a certain product or product category.

Referring back to Figure 8B, an Expired Decision Agent Manager function 154 maintains a list of Decision Agents 14 that are expired, that is, Decision Agents 14 that have completed their searches, whether successfully or not. The expired Decision Agents themselves are archived under the control of the consumer Personal Agent 12 that created them. Expired Decision Agent Manager 154 maintains indexes on the expired agents for quick searching by Producer's Demand Agents 16 that are, for example, calculating historical demand for a product.

Referring again to Figure 8A, an Active Demand Agent Manager function 132 maintains a list of all Demand Agents 16 that are currently calculating demand in this Market 18.

A Template Dispenser function 134 retrieves the Product Template 174 for a particular product. Product Template 174 describes the data that is available within the system about the particular product. Personal Agents 12 or 13 use the Template Dispenser 134 when consumers or providers are constructing ads or product search queries. Template Dispenser 134 consults the Product Template Manager 170 in a Product Database 32 (described in conjunction with Figure 9A) to collect the template data.

A Query Logger function 136 archives summary information about queries performed by Basic Search Engine 158 or Incremental Search Engine 162, so that historical data about queries may be quickly accessed without having to access the detailed data archived by the searching agent. For example, summary information about queries launched by Decision Agents 14 are logged so that Demand Agents 16 can perform routine demand calculations without having to access the archived Decision Agents 14.

-48-

An Historical Demand Search Engine 138 matches demand queries of Demand Agents 16 against the queries that have been previously logged by Query Logger 136, to identify Decision Agents 14 that previously, during a specified datetime range, had searched for a certain product or product category.

5 A Remote Database Adaptor 140 provides communication and session management services to connect to a database (a "remote database", not shown) belonging to a manufacturer or a provider. Remote Database Adaptor 140 also provides translation services to translate between the data formats used by a remote database and the data formats used by PDB 32. Remote Database Adaptor
10 140 allows a provider to submit ads directly from the provider's remote database into Market 18. Remote Database Adaptor 140 also allows access "by reference" to advertisement data that remains stored in a remote database; that is, the data is not copied into Agent System 10, but is accessed as needed. Market 18 includes a Remote Database Adaptor 140 for each provider that chooses to supply ads in this
15 manner; alternatively, a provider uses various functional components accessed via provider's Personal Agent 13 to place ads manually.

An Authorization function 142 restricts the placement and searching of ads in the Market 18 to authorized users only. General Markets 18a allow any authorized user of the system to place and search ads. A Restricted Market 18b
20 allows market access only to certain authorized users of the system. For example, a Restricted Market might be used by wholesalers marketing exclusively to distributors.

Product Database

25

Referring to Figure 9A, a Product Database 32 (PDB) comprises functional components:

-49-

a Database Administration function 166,
a Product Data Storage function 168,
a Product Template Manager function 170,
and, (optionally) some number of Remote Database Adaptors 172.

5 PDB 32 maintains generic data about products, to be referenced by ads placed by providers. Although PDB 32 is illustrated here as a single database (with several internal components) for ease of understanding, the contemplated PDB 32 will be split across several processors 38, as illustrated previously in Figure 3A.

Referring to Figure 9A, a Database Administration function 166 provides
10 conventional add, delete, update, query, and backup access for a System Administrator user to the other components of PDB 32.

A Product Data Storage function 168 stores data about different products, for example, product name, product model number, manufacturer's suggested retail price for product, etc.

15 A Product Template Manager function 170 maintains a set of Product Templates 174, one for each product listed in PDB 32. Product Template 174 describes the kinds of data that is kept in PDB 32 for a product. PDB 32 makes Product Templates 174 available to other components, for example, a Template Dispenser 134 as illustrated in Figure 8A. Referring now to Figure 9B, Product
20 Template 174 is comprised of a number of Product Template Entries 176 and, optionally, some Instructions for Use 178. Instructions for Use 178, if any, may be presented to the user when the user is supplying values for Product Template 174, to help the user decide what values to choose. A Product Template Entry 176 describes one property of a product. Referring to Figure 9C, Product Template
25 Entry 176 comprises several data components. A Keyword 180 names the property, for example, "Model Number" or "Brand Name". A Specification 182 indicates how values may be specified for the property, for example, as an integer

-50-

number, as a text string, as one item from an enumeration, etc. A set of Use Flags 184 indicate where the property is used, for example, when entering an ad, or when composing a search query. Product Template Manager 170 consults Use Flags 184 when collecting a set of Product Template Entries 176 to satisfy a
5 request from Template Dispenser 134. Referring to Figure 9D, there is shown an example of some Product Template Entries 176 for a hypothetical television set product, including sample values that would be stored in PDB 32. As illustrated, some products will contain advocate information. Some advocate information may be a ranking done by an independent rating organization. Another kind of advocate
10 information is an endorsement by a person or organization.

Referring again to Figure 9A, a Remote Database Adaptor 172 provides communication and session management services to connect to a product database (a "remote database", not shown) belonging to a manufacturer or a provider. Remote Database Adaptor 172 also provides translation services to translate
15 between the data formats used by a remote database and the data formats used by PDB 32. Remote Database Adaptor 172 is used to provide product data in real-time for manufacturers or providers that choose not to maintain product data directly in PDB 32 of Agent System 10, or to periodically update product data that is maintained directly in PDB 32. PDB 32 includes a Remote Database Adaptor
20 172 for each manufacturer or provider that chooses to supply product data in this manner; alternatively, a System Administrator may use Database Administration function 166 to maintain the data based on instructions from a manufacturer or provider.

25

Advertisement (Ad)

-51-

An advertisement (ad) is an offer to sell or buy a product. A placer is the user (person or organization) who is selling or buying.

Referring to Figure 10, a preferred embodiment of Ad 186 comprises various data components:

- 5 a Unique ID 188,
- a Buy/Sell Flag 190,
- a Reference to Placer component 192,
- a Reference to Market component 194,
- a Reference to Product Listing component 196,
- 10 a Product Template Value component 198,
- a Description component 200,
- a Price component 202,
- a Start Datetime component 204,
- and an Expiry Datetime component 206.

15 A Unique ID component 188 uniquely identifies this advertisement within Agent System 10. Unique ID 188 is generated automatically when the advertisement is created, and is never reused to identify a different ad, even after the advertisement is expired.

20 A Buy/Sell Flag 190 indicates that this advertisement is either an offer to buy or an offer to sell a product.

 A Reference to Placer component 192 identifies the provider Personal Agent 13 of the user placing Ad 186.

 A Reference to Market component 194 identifies a Market 18 in which Ad 186 is placed.

25 A Reference to Product Listing component 196 refers to standard data about the product in Product Listing 124 of Market 18. Brand name,

-52-

manufacturer, manufacturer's suggested retail price, etc. are examples of standard data.

A Product Template Values component 198 holds values, corresponding to the Product Template 174, that the placer specified when composing the advertisement. Offering price and locations of stores that carry the product are
5 examples of values.

A Description component 200 holds additional data, not in Product Listing 124 nor in Product Template Values 198, that the placer wishes to make known about the product.

10 A Price component 202 states the price at which the product is offered (for selling) or requested (for buying). Price 202 may also be a price range, especially for buy ads.

A Start Datetime component 204 states the date and time at which the advertisement becomes effective, that is, the point at which the advertisement will
15 be visible to Decision Agents 14 that are searching the market for products.

An Expiry Datetime component 206 states the date and time at which the advertisement expires, that is, the point at which the advertisement will no longer be visible to Decision Agents 14 that are searching the market for products. However, even after expiry, the advertisement is accessible through an Ad
20 Archive 96.

Flow Charts

The flow charts indicate the functional component primarily responsible for
25 carrying out a given task. For example, "Decision Composer" listed in the top part of a step means that the recited task is carried out by Decision Composer function

-53-

74 (see Figure 11). The flow charts also use subroutines to make it easier to follow high-level flows.

All communication between a user (either a consumer user or a provider user) and Agent System 10 is mediated by user's Personal Agent 12 or 13. The
5 flow charts and descriptions sometimes illustrate or state that a component receives input from the user or directs output to the user. These statements should be understood to mean that the component carries out the communication with the help of user's Personal Agent 12 or 13.

Some Figures illustrate example screen layout for input and output using a
10 Web browser interface. Underlined text represents hyperlinks.

Simple, conventional processes are not illustrated by flow charts. For example, the process whereby a user invokes Ad Tracker 94 to view the status of active ads placed by the user is not shown, since this type of process is well-known to practitioners.

15

Product Search

One of the major consumer uses of Agent System 10 is to assist a consumer in locating information about a product that is advertised for sale. It need not be
20 possible for the consumer to carry out the actual purchase within Agent System 10; it is only necessary that products be advertised within the system. However, when secure electronic transactions are available, it is anticipated that consumers will make actual purchases through the system. The product search process, while directly helping consumers, also generates consumer market data that is so crucial
25 to providers. This data tells providers the products for which consumers are searching, the criteria that are important to those consumers, and how many consumers are searching the various markets.

-54-

Referring to Figure 11, an overall method according to the present invention for searching for a product is referred to generally by reference numeral 220. When a user acting as a consumer decides to search within Agent System 10 for a product or product category, the consumer establishes a communications session with consumer's Personal Agent 12 (steps 222 - 224). Typically the consumer, using a personal computer, connects to consumer's Internet access provider, directs consumer's Web browser software to Agent System's electronic address (known as a URL), and enters a login name and password. A sample login screen is illustrated in Figure 39.

Referring again to Figure 11, the consumer specifies that a product search is desired, which invokes Decision Agent Manager 60 (step 226) to supervise the subsequent steps. A Compose Decision Query subroutine of Decision Composer 74 is called to assist the consumer in composing the query for the desired product (step 228).

Referring now to Figure 12A, there is shown a method for a Compose Decision Query subroutine, referred to generally by reference numeral 228. If the consumer wishes to specify a search that is similar to a previously performed search, Decision Agent Archive 80 displays a list of search queries from expired Decision Agents from which the consumer may select (steps 242 - 246). Decision Composer retrieves the Product Template mentioned in the selected Decision Agent, and also the current instructions, from the Market mentioned in the selected Decision Agent (step 248). The search criteria (values) from the expired Decision Agent are used to initialize the new search criteria (step 250).

Still referring to Figure 12A, if the consumer alternatively wishes to specify an entirely new search, the consumer selects a Market 18 in which to search (step 252). If the selected Market is a Restricted Market for which the consumer is not authorized, an error message is displayed to the consumer, and the

-55-

consumer is returned to the initial menu where another action may be selected (steps 254 - 260). If the consumer is authorized for the restricted Market, or if the Market is not restricted, the consumer selects a product for which to search (step 262). Decision Composer 74 retrieves Product Template and instructions from the Market's Template Dispenser 134 (step 264).

At this point Decision Composer 74 arranges to format and display the Product Template and the instructions (step 266). The consumer, following the instructions, completes the search criteria in the Product Template (step 268). When the consumer's interface is a Web browser, the Product Template is typically displayed as a combination of fill-in fields, selection lists, radio buttons, etc. as illustrated in the sample screens of Figure 40 (searching for consumer electronics) and Figure 41 (searching for automobiles).

Referring to Figure 12B, if the consumer wishes to perform an extended search, that is, a search that will continue for a period of time, the consumer enters a period of time for the search to continue (step 272). The extended search continues even when the consumer is not "on-line", that is, even when the consumer is not participating in a communication session with the Agent System. Search results are collected, as described below, for later delivery to the consumer. Alternatively, the consumer may instead choose an immediate search, that is, a search that will return results as soon as possible (step 274). The consumer need not be on-line to receive results from an immediate search; the results may be delivered later. The consumer may select a delivery media (e-mail, Web page display, etc.) and a delivery time and period (e.g., 6:00 p.m. daily, Monday noon weekly, etc.), or default media and time is noted (steps 276 - 280). At this point the Decision Query composition is complete (step 282).

Referring briefly to Figure 11, Decision Agent Factory 76 invokes a Create Decision Agent subroutine to create a new Decision Agent 14. Referring now to

-56-

Figure 13, a Create Decision Agent subroutine is referred to generally by reference numeral 230. Decision Agent Factory 76 creates a new Decision Agent 14 possessing a unique identifier (step 288). This unique identifier can be used to send messages to the agent even after the agent has expired (completed its task).

5 Decision Agent Factory 76 also initializes the other data components of the new agent by storing a reference to the Personal Agent of the consumer, a reference to the Market that is to be searched, the search expiry time, the delivery media, time, and period, and the query (search criteria) (step 290). Decision Agent Factory 76 logs the creation of the new agent with the new agent's Log function (step 292).
10 Now the new Decision Agent 14 is ready to be launched.

Referring again to Figure 11, Decision Agent Factory 76 delivers the new Decision Agent to the specified Market (step 232), where an Accept New Decision Agent subroutine is invoked (step 234). Referring now to Figure 14, an Accept New Decision Agent subroutine is referred to generally by reference numeral 234.
15 Active Decision Agent Manager 152 of the Sell Ad Manager 128 or Buy Ad Manager 130, as appropriate, accepts the new agent, logs the query from the agent to the Market's Query Logger function, and adds the agent to a queue of Immediate Agents 156 (steps 298 - 302).

Referring back to Figure 11, Decision Agent 14 is now ready to perform
20 the search for a product according to the consumer's criteria, so it invokes a Perform Decision Search subroutine (step 236). Referring to Figure 15, a Perform Decision Search subroutine is referred to generally by reference numeral 236. Decision Agent 14 performs an immediate search by invoking a Perform Immediate Search subroutine (step 308). Even if an extended search is chosen by
25 the consumer, an immediate search is done first to get initial results. Because the data to be searched resides within Agent System 10 or is easily and directly

-57-

accessible to Agent System 10, the search process can be faster than a search that employs agents which visit multiple Web sites or databases over the Internet.

Referring now to Figure 16, a Perform Immediate Search subroutine is referred to generally by reference numeral 308. An immediate search begins when Immediate Agents selects the next Decision Agent from its internal queue (step 320). Immediate Agents delivers the query from the chosen Decision Agent to Basic Search Engine 158 (step 322). Basic Search Engine 158 uses conventional database techniques to match the query against the ads in Active Ads, noting the ads that satisfy the query (step 324). The Decision Agent's Response Manager 108 collects references (step 326) to the matching ads found by Basic Search Engine. The Response Manager also sends a response to the Personal Agent that placed the advertisement (if the placer so desired and marked in the ad), providing real-time feedback to the placer. Immediate Agents then removes the Decision Agent from its internal queue and gives the Decision Agent back to Active Decision Agent Manager 152 (step 328).

Referring back to Figure 15, if an extended search was chosen by the consumer, Active Decision Agent Manager delivers the Decision Agent to Pending Agents (step 312), so that the query of the Decision Agent will continue to be matched against incoming ads until the Decision Agent's expiry time is reached. If the consumer chose only an immediate search, an End Decision Search subroutine is called to end the search (step 314).

Referring to Figure 17, an End Decision Search subroutine is referred to generally by reference numeral 314. The Decision Agent is removed from the queue of Immediate Agents (if the agent was performing an immediate search) or Pending Agents (if the agent was performing an extended search), and is delivered back to Personal Agent's Decision Agent Manager (step 334). Decision Agent Tracker 78 logs the end of the search (step 336). Decision Agent's Response

-58-

Manager 108 may still have undelivered search results. The results are held by the Response Manager until the consumer's specified delivery time arrives.

Referring now to Figure 18, an Extended Search subroutine is generally referred to by reference numeral 340. An episode of extended search begins each time a new advertisement arrives in a Market 18 (step 342). If there are no Decision Agents in Pending Agents (steps 344 - 346), there is no work to be done. However, if there are one or more Decision Agents pending for an extended search, the next agent is selected from the queue (step 348). If the agent's expiry time has been reached since the last sweep through the queue, the agent's search is stopped (step 352) with the End Decision Search subroutine previously illustrated. There is also a process (not shown) that periodically sweeps the queue and ends agents' searches, in case there is not enough advertisement activity in this market to activate Pending Agents on a regular basis. If the selected agent is not expired, Incremental Search Engine matches the agent's query against the data in the new advertisement (step 354). If the advertisement satisfies the query, Response Manager includes this advertisement in its list of results (steps 356 - 358), and notifies (if the advertisement so indicates) the placer of the advertisement that the advertisement was selected. This search process is repeated for each agent in Pending Agents.

Referring now to Figure 19, a Deliver Search Results subroutine is referred to generally by reference numeral 360. Immediate search results are delivered to the consumer when the consumer's desired delivery time is reached (which may be immediately if the consumer has so requested). Intermediate results from extended searched are delivered periodically according to the consumer's desired delivery period. When the desired delivery time is reached (step 362), Preference Manager 54 organizes the not-yet-delivered results according to the consumer's preferences (step 364). For example, results that mention favored brands are ordered before

-59-

results with less favored brands. Delivery Manager 56 formats the responses according to the consumer's desired delivery media (step 366). For example, if the consumer's desired delivery media is the Web, a Web page in HTML is generated. For another example, if the consumer desires e-mail delivery, a suitable
5 representation is generated. When formatting is complete, Delivery Manager 56 arranges the actual delivery of the search results (step 368). If the Decision Agent has completed its search, no more results will be forthcoming, so a subroutine Expire Decision Agent expires the Decision Agent (steps 370 - 372).

Referring to Figure 20, an Expire Decision Agent subroutine is referred to
10 generally by reference numeral 372. When a Decision Agent is expired, Expired Decision Agent Manager 154 logs the expired Decision Agent (step 378) so that Demand Agents can easily search through the expired Decision Agents of this market when calculating historical demand. The Decision Agent Tracker 78 notes that the agent is now expired (step 380), and the agent is permanently archived in
15 Decision Agent Archive 80 (step 382).

Quantify Demand

Quantifying demand is a major activity of Agent System 10. Demand is a
20 measure of the number of consumers interested in purchasing a product or interested in products in a category. Providers may quantify current demand or historical demand. Current demand measures the count of consumers that are currently searching for a product or searching within a product category. Historical demand measures the count of consumers that have searched for a
25 product, or searched within a product category, during a previous time period. Agent System 10 can not only calculate demand, it can also deliver a means of contacting those consumers (without revealing the actual identity of those

-60-

consumers). The demand search uses data that is generated by consumers as they search for products in the system.

Referring to Figure 21, an overall method according to the present invention for quantifying demand is referred to generally by reference numeral 386. When a user acting as a provider decides to quantify demand (perform a demand search) within Agent System 10 for a product or product category, the provider establishes a communication session with provider's Personal Agent 13 (steps 388 - 390). Typically the provider, using a personal computer, connects to provider's Internet access provider, directs provider's Web browser software to Agent System's electronic address (known as a URL), and enters a login name and password. A sample login screen has already been illustrated in Figure 39.

Referring to still to Figure 21, the provider specifies that a quantify demand function is desired, which invokes Demand Agent Manager 62 (step 392) to supervise the subsequent steps. A Compose Demand Query subroutine of Demand Composer 82 is called to assist the provider in composing the query that will gather the demand data (step 394).

Referring now to Figure 22, a Compose Demand Query subroutine is referred to generally by reference numeral 394. If the provider wishes to specify a demand search that is similar to a previously performed search, Demand Agent Archive 88 displays a list of search queries from expired Demand Agents from which the provider may select (steps 408 - 412). Demand Composer 82 retrieves the Product Template mentioned in the selected Demand Agent, and also the current instructions, from the Market mentioned in the selected Demand Agent (step 414). The search criteria (values) from the expired Demand Agent are used to initialize the new search criteria (step 416).

Still referring to Figure 22A, if the provider alternatively wishes to specify an entirely new search, the provider selects a Market 18 in which to search (step

-61-

418). If the selected Market is a Restricted Market for which the provider is not authorized, an error message is displayed to the provider, and the provider is returned to the initial menu where another action may be selected (steps 422 - 426). If the provider is authorized for the restricted Market, or if the Market is not restricted, the provider selects a product for which to search (steps 428). Demand Composer 82 retrieves Product Template and instructions from the Market's Template Dispenser 134 (step 430).

At this point Demand Composer 82 arranges to format and display the Product Template and the instructions (step 432). The provider, following the instructions, completes the search criteria in the Product Template (step 434). When the provider's interface is a Web browser, the Product Template is typically displayed as a combination of fill-in fields, selection lists, radio buttons, etc. For example, the provider might use screens similar to the sample consumer screens previously illustrated in Figures 40 and 41.

Referring to Figure 22B, the provider selects the type of demand to quantify (step 436). If the provider chooses to quantify current demand, Demand Composer fills in the datetime range to indicate that only currently active Decision Agents should be searched (steps 438 - 440). If the provider chooses to quantify historical demand, the provider selects a datetime range (steps 442 - 444) to indicate that only Decision Agents that were active during that datetime range should be searched. The provider may select a delivery media (e-mail, Web page display, etc.) and a delivery time and period (e.g., 6:00 p.m. daily, Monday noon weekly, etc.), or default media and time is noted (steps 446 - 450). At this point the Demand Query composition is complete (step 452).

Referring briefly to Figure 21, Demand Agent Factory 84 invokes a Create Demand Agent subroutine to create a new Demand Agent 16 (step 396). Referring now to Figure 23, a Create Demand Agent subroutine is referred to generally by

-62-

reference numeral 396. Demand Agent Factory 84 creates a new Demand Agent 16 possessing a unique identifier (step 458). This unique identifier can be used to send messages to the agent even after the agent has expired (completed its task). Demand Agent Factory 84 also initializes the other data components of the new agent by storing a reference to the Personal Agent of the provider, a reference to the Market that is to be searched, a datetime range indicating that only Decision Agents that were (are) active during that datetime range should be searched, the delivery media, time, and period, and the query (search criteria) (step 460).

5 Demand Agent Factory 84 logs the creation of the new agent with the new agent's Log function (step 462). Now the new Demand Agent 16 is ready to be launched.

Referring again to Figure 21, Demand Agent Factory 84 delivers the new Demand Agent to the specified Market (step 398), where an Accept New Demand Agent subroutine is invoked (step 400). Referring now to Figure 24, an Accept New Demand Agent subroutine is referred to generally by reference numeral 400. Active Demand Agent Manager 132 of the Market accepts the new agent, and logs the query from the agent to the Market's Query Logger 136 function (steps 468 - 470).

Referring back to Figure 21, Demand Agent 16 is now ready to perform the search for Decision Agents 14 that satisfy the provider's criteria, so it invokes a Perform Demand Search subroutine (step 402). Referring to Figure 25, a Perform Demand Search subroutine is referred to generally by reference numeral 402. Demand Agent 16 determines if the search is for current or historical demand, and invokes an appropriate subroutine, either a Perform Current Demand subroutine, or a Perform Historical Demand subroutine (steps 476 - 480).

Referring now to Figure 26, a Perform Current Demand subroutine is referred to generally by reference numeral 478. Demand Agent 16 delivers its query to a Current Demand Search Engine 164 (step 490). Current Demand

-63-

Search Engine 164 matches the query, using conventional database techniques, against the Decision Agents that are listed in Immediate Agents 156 and Pending Agents 160 (step 492), as these Decision Agents are the ones that are currently searching for products. During the search, Current Demand Search Engine 164
5 collects references to Decision Agents 14 whose queries satisfy the query of the Demand Agent 16. For example, if the Demand Agent is looking for consumers who are currently looking for sports shoes, the Current Demand Search Engine will collect references to Decision Agents that are searching for sports shoes. Current Demand Search Engine 164 delivers the collected list of references to the
10 Demand Agent (step 494). When the search is complete, Demand Agent notifies Active Demand Agent Manager 132 that the search is complete (step 496).

Referring now to Figure 27, a Perform Historical Demand subroutine is referred to generally by reference numeral 480. Demand Agent 16 delivers its query to an Historical Demand Search Engine 138 (step 502). Historical Demand
15 Search Engine 138 matches the query, using conventional database techniques, against the expired queries that are kept in Query Logger 136 (step 504). During the search, Historical Demand Search Engine 138 collects references to expired queries, and the Decision Agents to which they belong, that were active during the specified datetime range of the Demand Agent's query and that otherwise satisfy
20 the Demand Agent's query. Historical Demand Search Engine 138 delivers the collected list of references to the Demand Agent (step 506). When the search is complete, Demand Agent notifies Active Demand Agent Manager 132 that the search is complete (step 508).

Referring back to Figure 25, the search is complete, so Active Demand
25 Agent Manager 132 delivers the Demand Agent back to Personal Agent's Demand Agent Manager 62 (step 482), and Personal Agent's Demand Agent Tracker 86 logs the search completion (step 484). The demand search results are held by the

-64-

Demand Agent 16 until the provider's desired delivery time (which may be immediately if the provider has so requested).

The search has now ended. It remains for Demand Agent 16 to deliver the demand search results to the provider. Referring to Figure 28, a Deliver Demand
5 Results subroutine is referred to generally by reference numeral 512. When the desired delivery time is reached (step 514), Preference Manager 54 organizes the not-yet-delivered results according to the provider's preferences (step 516). For example, the provider may prefer to see only numeric totals, or the provider may prefer to see a detailed listing of all the Decision Agent queries that satisfied the
10 demand search. Delivery Manager 56 formats the responses according to the provider's desired delivery media (step 518). For example, if the provider's desired delivery media is the Web, a Web page in HTML is generated. If the provider desires e-mail delivery, a suitable representation is generated. When formatting is complete, Delivery Manager 56 arranges the actual delivery of the
15 search results (step 520). The Demand Agent has completed its task, so it can be expired (step 522).

Referring to Figure 29, an Expire Demand Agent subroutine is generally referred to by reference numeral 522. The Demand Agent Tracker 86 notes that the agent is now expired (step 528), and the agent is permanently archived in
20 Demand Agent Archive 88 (step 530).

Place Ad

Both providers and consumers may place ads in Agent System 10. An
25 advertisement may be an offer to sell or an offer to buy. A placed advertisement becomes effective at a particular time and expires at a particular time, and searching Decision Agents consider an advertisement only during the ad's effective

-65-

time. Even after ads expire, they may be accessed for historical reasons. Users may place ads manually, or they may cause ads to be loaded from or referenced from remote systems via a remote database adapter. Referring to Figure 30, an overall method according to the present invention for placing a sell or buy advertisement is referred to generally by reference numeral 534. When a provider (a user acting in the role of a provider) desires to place an advertisement manually within Agent System 10 for a product, the provider establishes a communication session with provider's Personal Agent 13 (steps 536 - 538). Typically the provider, using a personal computer, connects to provider's Internet access provider, directs provider's Web browser software to Agent System's electronic address (known as a URL), and enters a login name and password. A sample login screen has already been illustrated in Figure 39.

Referring to Figure 30, the provider invokes Sell Ad Manager 128 or Buy Ad Manager 130 as appropriate (step 540) to supervise the subsequent steps of placing the advertisement. A Compose Ad subroutine of Ad Composer 90 is called to assist the provider in composing the advertisement (step 542).

Referring now to Figure 31, a Compose Ad subroutine is referred to generally by reference numeral 542. The provider selects the type of ad: a sell advertisement (an offer to sell) or a buy advertisement (an offer to buy) (step 552). If the provider wishes to compose an advertisement that is similar to a previously placed ad, Ad Archive 96 displays a list of expired ads from Ad Archive 96 from which the provider may select (steps 554 - 558). Ad Composer 90 retrieves the Product Template mentioned in the selected ad, and also the current instructions, from the Market mentioned in the selected advertisement (step 560). The values from the selected advertisement are used to initialize the new advertisement (step 562).

-56-

Still referring to Figure 31A, if the provider alternatively wishes to specify an entirely new ad, the provider selects a Market 18 in which the advertisement will be placed (step 564). If the selected Market is a Restricted Market for which the provider is not authorized, an error message is displayed to the provider, and the provider is returned to the initial menu where another action may be selected (steps 566 - 570). If the provider is authorized for the restricted Market, or if the Market is not restricted, the provider selects a product for which to search (steps 568 - 574). Ad Composer 90 retrieves Product Template and instructions from the Market's Template Dispenser 134 (step 576). If the particular product is not listed in the Market, the provider instead indicates a "generic" ad, and Template Dispenser 134 supplies a generic template that can be used for any product suitable for the Market.

Ad Composer 90 arranges to format and display the Product Template and the instructions (step 578). The provider, following the instructions, selects and enters values describing the product in the Product Template (step 580), adding additional description if desired. When the provider's interface is a Web browser, the Product Template is typically displayed as a combination of fill-in fields, selection lists, radio buttons, etc. as illustrated in the sample screen of Figure 42.

Referring again to Figure 31B, the provider enters the price for the product (step 582). Generally, a sell advertisement will contain a specific price for the product, while a buy advertisement will contain a price range. The provider specifies the datetime that the advertisement should become effective and the datetime that the advertisement should expire (step 584). This allows providers to compose batches of ads ahead of time, for example with lower prices during a sale, and arrange for the ads to become affective when the sale starts. For receiving responses to the ad, the provider may select a delivery media (e-mail, Web page display, etc.) and a delivery time and period (e.g., immediately, 6:00

-67-

p.m. daily, Monday noon weekly, etc.), or default media and time is noted (steps 586 - 590). At this point the advertisement composition is complete (step 592), and Ad Composer 90 can create the actual advertisement with a Create Ad subroutine (step 594).

5 Referring now to Figure 32, a Create Ad subroutine is referred to generally by reference numeral 594. Ad Composer 90 creates a new Ad 186 possessing a unique identifier (step 600). This unique identifier can be used to access Ad 186 even after the advertisement has expired. Ad Composer 90 indicates that this is a buy advertisement or sell advertisement. Ad Composer 90 also inserts values (step
10 602) for the other data components as specified by the provider in previous steps: Template Values, Description, Price, Start and Expiry Datetime. Ad Composer 90 inserts references to the Personal Agent 12 or 13 that is controlling the advertisement creation, and to the Personal Agent 12 or 13 of the principal (buyer or seller). Ad Composer 90 notes in which Market 18 the advertisement is to be
15 placed, and inserts a reference to the standard data about the product from the Product Listing 124. Ad Tracker 94 logs the creation of the new Ad 186 (step 604). Now the new Ad is ready to be delivered.

Referring again briefly to Figure 30, Ad Delivery function 92 delivers the newly created Ad 186 to the Sell Ad Manager 128 or Buy Ad Manager 130, as
20 appropriate, of the provider's chosen Market (step 544), and the Ad Manager invokes an Accept New Ad subroutine to incorporate the new advertisement (step 546).

Referring now to Figure 33, an Accept New Ad subroutine is referred to generally by reference numeral 546. Sell Ad Manager 128 or Buy Ad Manager
25 130, as appropriate, accepts the Ad 186 and checks the ad's Start Datetime to see if the advertisement should become active (effective) now. (steps 612 - 614). If it is not time for the advertisement to become effective, the advertisement is passed

-68-

to Future Ads 148, which holds the advertisement until the appropriate future time (step 616). Future Ads arranges to give the advertisement back to the appropriate Ad Manager for activation at the appropriate time.

5 If it is time for the advertisement to become active, Active Ads 146 inserts the advertisement into its internal queue and Ad Indexing 144 indexes the advertisement for searching (steps 618 - 620). Active Ads notifies Active Decision Agent Manager 152 that a new advertisement has arrived (step 622), so that ongoing extended searches may be matched against the new advertisement.

10 Once Ad 186 is in Active Ads 146, the advertisement is available for searching by Decision Agents 14 that are looking for products. The advertisement remains available for searching until its Expiry Datetime is reached, when Ad Manager invokes an Expire Ad subroutine.

15 Referring now to Figure 34, an Expire Ad subroutine is referred to generally by reference numeral 626. Sell Ad Manager 128 or Buy Ad Manager 130, as appropriate, removes the Ad 186 from Active Ads 146 so that the advertisement is no longer visible to searching Decision Agents 14 (step 628). The indexes for the advertisement are transferred to Expired Ads 150 (step 630) to make historical searches easier. Ad Tracker 94 logs the expiration of the advertisement (step 632). Ad Archive 96 permanently archives the advertisement (step 634). Even though the advertisement is expired, it can still be referenced out
20 of the Ad Archive for historical searches.

Target Consumers

25 To "target" a message is to select message recipients according to certain criteria intended to yield recipients that are interested in receiving the information, as opposed to delivering the message to a wider audience where fewer recipients

-69-

are truly interested in the information. Within Agent System 10, targeting consumers is a process of targeting Personal Agents 12 that represent consumers that satisfy the desired criteria. Providers use the targeting process, for example, to send ads to consumers that have previously searched in a particular Market 18
5 or for a particular product. Providers may also use targeting to offer a product at a different price to a different groups of consumers. Providers might also send market surveys, brand name awareness notices, etc.

Providers can also target consumers who have recently searched a Market 18 to deliver a "reason for sale" or "lost sales" questionnaire. The questionnaire
10 inquires if the consumer actually bought a product, and if the purchased product was the one offered by the inquiring provider or some other provider. The questionnaire typically includes a list of sales reasons, that is, reasons why the consumer purchased the product or purchased from the inquiring provider. For example, some sales reasons are: price suitable, available in desired color, a
15 particular special feature, etc. The questionnaire also typically includes a list of lost sales reasons, that is, reasons why the consumer purchased a competing product or purchased from another provider. For example, some lost sales reasons are: price too high, prefer another brand, store location not convenient, etc. The questionnaire may also include a place for general comments from the consumer.
20 By analyzing returned questionnaires, the provider gains valuable information about why a sale was gained or why a competitor got the sale.

Referring to Figure 35, an overall method according to the present invention for targeting consumers is referred to generally by reference numeral 638. When a provider (a user acting in the role of a provider) desires to target a
25 message to a select group of consumers, the provider establishes a communication session with provider's Personal Agent 13 (steps 640 - 642). Typically the provider, using a personal computer, connects to provider's Internet access

-70-

provider, directs provider's Web browser software to Agent System's electronic address (known as a URL), and enters a login name and password. A sample login screen has already been illustrated in Figure 39.

5 Referring again to Figure 35, Target Manager 66 assists the provider in identifying the target set of Personal Agents that will receive the message by invoking a Target Personal Agents subroutine. Referring to Figure 36, a Target Personal Agents subroutine is referred to generally by reference numeral 644. The provider executes a process of refinement (steps 662 - 670) to collect references to appropriate Personal Agents.

10 The provider may start by quantifying demand, previously described in conjunction with Figure 21, and use the set of consumer Personal Agents 12 thereby identified. For example, the provider may quantify current demand for sports shoes, and collect the Personal Agents that currently have Decision Agents that are looking for sports shoes.

15 Alternatively, the provider may select, from Demand Agent Archive 88, a Demand Agent 16 that previously identified an appropriate set of consumer Personal Agents 12.

As yet another alternative, the provider may use Personal Agent Search Engine 26 to collect references to a set of consumer Personal Agents 12 that have 20 certain preferences or demographic characteristics. For example, the provider may search for Personal Agents that list a preference for a certain favorite brand, or for Personal Agents whose owners are males between the ages of 25 and 40. Personal Agent Search Engine 26 also identifies Personal Agents that are willing to accept unsolicited notices only if accompanied by a consideration. In this context, 25 Personal Agent Search Engine 26 automatically filters out Personal Agents that belong to persons or organizations that are no longer users of the system.

-71-

The provider may keep applying search criteria to filter the set of consumer Personal Agents 12 at will, until the provider is satisfied that an appropriate set has been identified. By choosing suitable search criteria, the provider may also select a set of consumers to quantify anticipated or future demand. For example, 5 the provider may select consumers that have recently searched a real estate market, and anticipate that those consumers will soon desire mortgage lending information.

Referring again to Figure 35, the provider composes the actual message to be sent (step 646). For example, if the provider wants to send an ad, the provider may use the Ad Composer 90 to assist in the composition, using a Compose Ad 10 subroutine previously described. Other notices, messages, consumer surveys, etc. could also be composed (not shown) for delivery.

The message is delivered to each consumer Personal Agent 12 that was identified with the help of Target Manager 66 (step 648). For example, if the message is an ad, Ad Delivery 92 arranges the delivery.

15 Consumers receive the messages via their Personal Agents 12 and compose replies if they wish (step 650). This may involve filling out an on-line form presented by the provider. The replies are sent back to the originating provider.

The provider may have specified a consideration amount to be paid to consumers that reply to the message (step 652). If so, Target Manager 66 arranges 20 to send a consideration notice to the consumer Personal Agent 12 of each consumer that replies (step 654). Consideration Account 67 of each consumer receiving a consideration notice credits the consideration account with the amount of the consideration (step 656).

25 **Reject Unsolicited Message**

-72-

Even unsolicited advertisements and other unsolicited messages that are rejected by a consumer's Personal Agent 12 become sources of market data to a provider, if the rejection generates a rejection reason back to the provider.

Referring to Figure 37, an overall method according to the present invention for rejecting an unsolicited message is referred to generally by reference numeral 674. When a consumer's Delivery Manager 56 receives an ad, market survey, notice, or other message that is unsolicited, Delivery Manager 56 matches the data in the message against the preferences maintained by Preference Manager 54 (steps 676 - 678). If the message content does not violate any of the consumer's preferences, the message is delivered by Delivery Manager 56 in the usual fashion according to the delivery media and delivery time preferences of the consumer (step 682).

If, however, the message content violates the consumer's preference in some way, Delivery Manager 56 composes a rejection message indicating the reason for rejection, and sends the rejection message back to the Personal Agent 13 of the provider that originated the unsolicited message (steps 684 - 686).

For example, if a provider sends, to Personal Agent 13, an unsolicited advertisement about sports shoes, specifying a consideration amount of fifty cents, and the user has previously specified a consideration preference of seventy-five cents, Delivery Manager 56 will reject the advertisement and reply with a rejection message indicating that the consideration amount must be at least seventy-five cents.

As another example, perhaps a provider sends an unsolicited advertisement about a Chinese food dinner, and the advertisement does not specify the MSG content of the food. If the user has specified a preference for "no MSG", Delivery Manager 56 will reject the advertisement and reply with a rejection message indicating that the consumer prefers food without MSG.

-73-

The provider has gained valuable market information about consumer preferences, even though the provider's message was not successfully delivered.

Simulate Demand

5

Providers may also simulate demand using current market data. To simulate demand is to determine the demand for a hypothetical product or service, or for an actual product or service with different features or pricing.

Referring to Figure 38A, when a provider decides to simulate demand for a product, the provider composes and places an Ad 186 in a Market 18 in the regular fashion, but the provider additionally marks the Ad as "invisible" (step 694). The provider composes the Ad to represent the hypothetical product, or the product with different features or pricing. The Ad is (during its effective datetime) available to be matched by consumers' Decision Agents 14 performing decision searches.

15

When a searching Decision Agent 14 matches the invisible Ad 186, the Ad is referred to the Response Manager 108 and the Preference Manager 54 ranks the Ad in the normal fashion (steps 696 - 698). However, because the Ad is marked invisible, the Response Manager additionally sends a reply to the provider indicating that the invisible Ad was matched by a Decision Agent 14, and indicating the ranking of the invisible Ad (step 700). The Ad 186 reference is not removed from the Decision Agent 14, but remains with the other Ad references of the Decision Agent in the normal fashion. When Delivery Manager 56 is preparing search results for delivery to the consumer, it does not include the invisible Ad (step 702), so that the consumer remains unaware that the invisible Ad existed. The provider has collected valuable market data without annoying the consumers who generated the data.

20
25

-74-

Referring to Figure 38B, a variation of this process allows providers to "replay" a product offering, that is, to estimate what effect an advertisement would have had if, for example, the price had been lower. This process variation matches the invisible advertisement against expired Decision Agents 14 over a
5 datetime range.

Other Embodiments

The foregoing description is of a preferred embodiment of the invention. Other embodiments are anticipated. For example, it is expected that future
10 embodiments of the invention will use a variety of communication devices, such as, but not limited to, facsimile machines, pagers, Personal Digital Assistants (PDAs), Network Computers (NCs), postal mail, telephone voice recognition, satellite links, video cable, etc.

It is also anticipated that, in the future, the system will further comprise
15 actual purchase transactions.

It is also anticipated that additional kinds of data will be collected by the system, and additional methods of analysis of such data will be developed.

Conclusion

20 From the foregoing it will be seen that this invention is well adapted to attain all of the ends and objectives hereinabove set forth, together with other advantages which are inherent to the apparatus.

It will be understood that certain features and sub-combinations are of utility and may be employed without reference to other features and
25 sub-combinations. This is contemplated by and is within the scope of the claims.

The foregoing description of the preferred embodiment of the invention has been presented for the purposes of illustration and description. It is not intended

-75-

to be exhaustive or to limit the invention to the precise form disclosed. Many modifications and variations are possible in light of the above teaching.

5

-76-

The Claims

1. A computer network agent system for providing communication between
an anonymous potential consumer of products which can be goods or
5 services and a provider of such products, comprising in combination:
a consumer personal agent for receiving product queries from the potential
consumer and transmitting product recommendations to the potential
consumer;
a decision agent for receiving anonymous product queries from the
10 consumer personal agent and transmitting product recommendations
to the consumer personal agent;
a provider personal agent for receiving demand queries from the provider
and transmitting quantified demand information to the provider;
a demand agent for receiving demand queries from the provider personal
15 agent and transmitting quantified demand information to the
provider personal agent; and
a market for gathering information from the agents, organizing the
information and distributing organized information to the agents.

- 20 2. A computer network agent system according to Claim 1 wherein the
consumer personal agent comprises, in combination:
a unique identifier function for maintaining an identifier that uniquely
identifies the consumer personal agent within the agent system;
an owner manager function for maintaining data about the consumer;
25 an individual firewall for protecting the data about the consumer from other
agents, controlling all access to data that is maintained by the

-77-

- various internal functions, ensuring that only authenticated and authorized agents and users can access private data;
- a preference manager for maintaining data about the preferences of the consumer;
- 5 a delivery manager for accepting messages generated by other agents or other components of the system, that are directed to the consumer and delivers those messages according to the consumer's desired delivery time and delivery media and rejecting unsolicited messages, unless an unsolicited message meets the consumer's
- 10 preferred criteria as maintained by the preference manager; and
- a decision agent manager for assisting the consumer with the creation and management of decision agent.
3. A computer network agent system according to Claim 1 wherein the
- 15 consumer personal agent further comprises a consideration account.
4. A computer network agent system according to Claim 2 wherein the decision agent manager comprises, in combination:
- a decision composer for assisting the consumer in composing queries to be
- 20 executed by decision agents;
- a decision agent factory for creating a new decision agent to carry out a query; and
- a decision agent tracker for enabling the consumer to monitor the activities of decision agents that have not completed their tasks.
- 25

-78-

5. A computer network agent system according to Claim 4 wherein the decision agent manager further comprises a decision agent archive for storing and accessing decision agents that have completed their tasks.
- 5 6. A computer network agent system according to Claim 1 wherein the provider personal agent comprises, in combination:
a unique identifier function for maintaining an identifier that uniquely
identifies the provider personal agent within the agent system;
an owner manager function for maintaining data about the provider;
10 an individual firewall for protecting the data about the provider from other agents, controlling all access to data that is maintained by the various internal functions, ensuring that only authenticated and authorized agents and users can access private data;
a preference manager for maintaining data about the preferences of the
15 provider;
a delivery manager for accepting messages generated by other agents or other components of the system, that are directed to the provider and delivers those messages according to the provider's desired delivery time and delivery media and rejecting unsolicited
20 messages, unless an unsolicited message meets the provider's preferred criteria as maintained by the preference manager; and
a demand agent manager for assisting the provider with the creation and management of demand agents.
- 25 7. A computer network agent system according to Claim 6 wherein the demand agent manager comprises, in combination:

-79-

- a demand composer for assisting the provider in composing queries to be executed by demand agents;
- a demand agent factory for creating a new demand agent to carry out a query; and
- 5 a demand agent tracker for enabling the provider to monitor the activities of demand agents that have not completed their tasks.
8. A computer network agent system according to Claim 7 wherein the demand agent manager further comprises a demand agent archive for
- 10 storing and accessing demand agents that have completed their tasks.
9. A computer network agent system according to Claim 6 further comprising an advertisement manager for assisting the provider in placing advertisements.
- 15
10. A computer network agent system according to Claim 9 wherein the advertisement manager comprises, in combination:
- an advertisement composer for retrieving product information;
- an advertisement delivery function for delivering an advertisement to a
- 20 desired destination;
- an advertisement tracker for monitoring the activity of the advertisement, including any messages received in response to the advertisement, until the advertisement expires or is canceled by the user.
- 25 11. A computer network agent system according to Claim 10 wherein the advertisement manager further comprises an advertisement archive for storing and accessing advertisements that are expired.

-80-

12. A computer network agent system according to Claim 6 wherein the provider personal agent further comprises a target manager for assisting the provider in identifying consumer personal agents to which targeted messages may be delivered.
- 5
13. A computer network agent system according to Claim 1 wherein the decision agent comprises, in combination:
- 10 a unique identifier function for maintaining an identifier that uniquely identifies this decision agent within the agent system;
- a personal agent reference for holding a copy of the unique identifier of the consumer personal agent that controls this decision agent;
- a market reference for indicating in which market the decision agent should search;
- 15 an expiry function for indicating how long the decision agent should continue searching;
- a query for describing the product or product category for which to search;
- a response manager for receiving search results and returning the search results to the consumer personal agent; and
- 20 a log function for storing records of the activities of the decision agent.
14. A computer network agent system according to Claim 1 wherein the demand agent comprises, in combination:
- 25 a unique identifier function for maintaining an identifier that uniquely identifies this demand agent within the agent system;
- a personal agent reference for holding a copy of the unique identifier of the provider personal agent that controls this demand agent;

-81-

- a market reference for indicating in which market or markets the demand agent should search;
- a datetime range function for indicating that demand should be quantified over the date/time range specified;
- 5 a demand query for describing a product or product category query that can be matched against the queries of decision agents; and
- a log function for storing records of the activities of the demand agent for later consultation by other components of the agent system.
- 10 15. A computer network agent system according to Claim 1 wherein the market comprises, in combination:
- a product listing function for maintaining a list of the products that can be advertised in this market;
- a cross reference manager for maintaining cross references to other markets
- 15 that carry similar products;
- a sell advertisement manager for accepting advertisements of offers to sell that are submitted by consumer personal agents;
- a buy advertisement manager for accepting advertisements of offers to buy that are submitted by consumer personal agents;
- 20 an active demand agent manager for maintaining a list of all demand agents that are currently calculating demand in this market;
- a template dispenser for retrieving data that is available within the agent system about a particular product;
- a query logger for archiving summary information about queries so that
- 25 historical data about queries may be quickly accessed without having to access the detailed data which has been archived; and

-82-

an historical demand search engine for matching demand queries of
demand agents against the queries that have been previously logged
by the query logger to identify decision agents that previously,
during a specified datetime range, had searched for a certain
product or product category.

5

16. A computer network agent system according to Claim 15 wherein the
market further comprises a remote database adaptor for providing
communication and session management services to connect to a remote
database belonging to a manufacturer or a provider and translating between
the data formats used by the remote database and the data formats used by
the market.

10

17. A computer network agent system according to Claim 15 wherein the
market further comprises an authorization function for restricting the
placement and searching of advertisements in the market to authorized
consumers and providers only.

15

18. A computer network agent system according to Claim 15 wherein the buy
advertisement manager further comprises, in combination:
an advertisement indexer for maintaining indexes for quick searching of the
advertisements by product and vendor characteristics;
an active advertisements function for maintaining information about
advertisements that are currently active, wherein an active decision
agent manager is notified as each new advertisement is accepted by
the active advertisements function so that pending searches can be
matched against the new advertisement;

20

25

-83-

a future advertisements function for maintaining advertisements that have been submitted to the market, but are not yet active because their effective datetime has not yet been reached;

an expired advertisements function for maintaining an archive of advertisements that have expired, for later analysis;

an active decision agent manager for maintaining a list of all decision agents that are currently searching this market for products; and

an expired decision agent manager for maintaining a list of decision agents that have completed their searches, whether successfully or not.

10

19. A computer network agent system according to Claim 18 wherein the active decision agent manager comprises, in combination:

an immediate agents function for keeping track of decision agents that are performing searches that are to be performed and results returned as soon as practical;

15

a basic search engine for executing queries associated with each incoming decision agent by consulting active advertisements to match the query against all active advertisements;

a pending agents function for keeping track of decision agents that are performing extended searches;

20

an incremental search engine for matching each new advertisement against the queries of the pending decision agents; and

a current demand search engine for matching demand queries of demand agents against the queries of decision agents that are residing in immediate agents or pending agents, to identify active decision agents that are searching for a certain product or product category.

25

-84-

20. A computer network agent system according to Claim 15 wherein the sell advertisement manager further comprises, in combination:
- an advertisement indexer for maintaining indexes for quick searching of the advertisements by product and vendor characteristics;
 - 5 an active advertisements function for maintaining information about advertisements that are currently active, wherein an active decision agent manager is notified as each new advertisement is accepted by the active advertisements function so that pending searches can be matched against the new advertisement;
 - 10 a future advertisements function for maintaining advertisements that have been submitted to the market, but are not yet active because their effective datetime has not yet been reached;
 - an expired advertisements function for maintaining an archive of advertisements that have expired, for later analysis;
 - 15 an active decision agent manager for maintaining a list of all decision agents that are currently searching this market for products; and
 - an expired decision agent manager for maintaining a list of decision agents that have completed their searches, whether successfully or not.
- 20 21. A computer network agent system according to Claim 20 wherein the active decision agent manager comprises, in combination:
- an immediate agents function for keeping track of decision agents that are performing searches that are to be performed and results returned as soon as practical;
 - 25 a basic search engine for executing queries associated with each incoming decision agent by consulting active advertisements to match the query against all active advertisements;

-85-

- a pending agents function for keeping track of decision agents that are performing extended searches;
- an incremental search engine for matching each new advertisement against the queries of the pending decision agents; and
- 5 a current demand search engine for matching demand queries of demand agents against the queries of decision agents that are residing in immediate agents or pending agents, to identify active decision agents that are searching for a certain product or product category.
- 10 22. A method for searching for a product by a consumer, comprising in combination the steps of:
- selecting a decision agent manager to supervise the subsequent steps;
- composing a decision query;
- creating a decision agent;
- 15 delivering the decision agent to a specified market;
- accepting the decision agent by the market; and
- searching for the product.
- 20 23. A method for searching for a product according to Claim 22 wherein the step of composing a decision query comprises, in combination, the steps of:
- if specifying an entirely new search, selecting a market in which to search;
- if the selected market is a restricted market for which the consumer is not authorized, displaying an error message to the consumer, and
- 25 returning to the initial menu where another action may be selected;
- if the consumer is authorized for the restricted market, or if the market is not restricted, selecting the product for which to search;

-86-

retrieving a product template and instructions from a market template
dispenser;

formatting and displaying the product template and the instructions;

completing the search criteria in the product template;

5 if performing a search that will continue for a period of time, entering a
period of time for the search to continue;

if performing a search that will return results as soon as possible, indicating
such; and

selecting a delivery media and a delivery time and period.

10

24. A method for searching for a product according to Claim 22 wherein the
step of creating a decision agent comprises, in combination, the steps of:
creating a decision agent with a unique identifier;

15

storing a reference to the personal agent of the consumer, a reference to the
market that is to be searched, the search expiry time, the delivery
media, time, and period, and the query; and

logging the creation of the decision agent with the new agent's log
function.

20

25. A method for searching for a product according to Claim 22 wherein the
step of accepting the decision agent comprises, in combination, the steps
of:

accepting the new agent by an active decision agent manager;

logging the query from the agent to the market's query logger function;

25

and

adding the agent to a queue of immediate agents.

-87-

26. A method for searching for a product according to Claim 22 wherein the step of searching for the product comprises, in combination, the steps of: searching for the product as soon as possible;
if performing a search that will continue for a period of time, delivering
5 the decision agent to queue of pending agents; and
if performing a search that will return results as soon as possible, ending the decision search.
27. A method for searching for a product according to Claim 26 wherein the
10 step of searching for the product as soon as possible comprises, in combination, the steps of:
selecting the next decision agent from the queue;
delivering the query from the decision agent to a search engine;
matching the query against active advertisements;
15 collecting matching advertisements;
responding to each placer of an advertisement to indicate that the advertisement was collected; and
giving the decision agent back to the active decision agent manager.
- 20 28. A method for quantifying demand by a provider for a product, comprising in combination the steps of:
selecting a demand agent manager to supervise the subsequent steps;
composing a demand query;
creating a demand agent;
25 delivering the demand agent to a market;
accepting the demand agent by the market; and
searching demand.

-88-

29. A method for quantifying demand according to Claim 28, wherein the step of composing a demand query comprises, in combination, the steps of:
if specifying an entirely new search, selecting a market in which to search;
if the selected market is a restricted market for which the provider is not
5 authorized, displaying an error message to the provider, and
returning to the initial menu where another action may be selected;
if the provider is authorized for the restricted market, or if the market is
not restricted, selecting the product for which to search;
retrieving a product template and instructions from a market template
10 dispenser;
formatting and displaying the product template and the instructions;
completing the search criteria in the product template;
if performing a search for historical demand, entering a date and time
range;
15 if performing a search for current demand, indicating such; and
selecting a delivery media and a delivery time.
30. A method for quantifying demand according to Claim 28, wherein the step
of creating a demand agent comprises, in combination, the steps of:
20 creating a decision agent with a unique identifier;
storing a reference to the personal agent of the provider, a reference to the
market that is to be searched, the search expiry time, the delivery
media, time, and period, and the query; and
logging the creation of the new agent with the new agent's log function.
25
31. A method for quantifying demand according to Claim 28, wherein the step
of accepting the demand agent comprises, in combination, the steps of:

-89-

accepting the agent by an active demand agent manager; and
logging the query from the demand agent to the market's query logger
function.

- 5 32. A system for electronic commerce wherein market data can be collected
and analyzed comprising:
a plurality of consumer agents, each associated with and capable of
communicating with a consumer;
a consumer data base comprising consumer preference data associated with
10 each said consumer;
a plurality of provider agents, each associated with and capable of
communicating with a provider; and
a data base of offers to sell and offers to buy;
wherein said consumer agents conceal the identity of the associated
15 consumer from agents not associated with said consumer in said
system;
wherein said consumer agents are capable of searching said data base of
offers to sell and offers to buy;
wherein said consumer agent searching generates persistent market data;
20 wherein said provider agents are capable of searching said consumer
database and said persistent market data; and
wherein said provider agents analyze said consumer database and said
persistent market data to quantify consumer demand.
- 25 33. The system of claim 32 wherein said consumer data base furthers
comprises non-identifying demographic data.

-90-

34. A computer network agent system according to Claim 32 wherein each consumer personal agent comprises, in combination:
- a unique identifier function for maintaining an identifier that uniquely identifies the consumer personal agent within the agent system;
 - 5 an owner manager function for maintaining data about the consumer;
 - an individual firewall for protecting the data about the consumer from other agents, controlling all access to data that is maintained by the various internal functions, ensuring that only authenticated and authorized agents and users can access private data;
 - 10 a preference manager for maintaining data about the preferences of the consumer;
 - a delivery manager for accepting messages generated by other agents or other components of the system, that are directed to the consumer and delivers those messages according to the consumer's desired
 - 15 delivery time and delivery media and rejecting unsolicited messages, unless an unsolicited message meets the consumer's preferred criteria as maintained by the preference manager; and
 - a decision agent manager for assisting the consumer with the creation and management of decision agent.
 - 20
35. A computer network agent system according to Claim 32 wherein each provider personal agent comprises, in combination:
- a unique identifier function for maintaining an identifier that uniquely identifies the provider personal agent within the agent system;
 - 25 an owner manager function for maintaining data about the provider;
 - an individual firewall for protecting the data about the provider from other agents, controlling all access to data that is maintained by the

-91-

- various internal functions, ensuring that only authenticated and authorized agents and users can access private data;
- a preference manager for maintaining data about the preferences of the provider;
- 5 a delivery manager for accepting messages generated by other agents or other components of the system, that are directed to the provider and delivers those messages according to the provider's desired delivery time and delivery media and rejecting unsolicited messages, unless an unsolicited message meets the provider's
- 10 preferred criteria as maintained by the preference manager; and
- a demand agent manager for assisting the provider with the creation and management of demand agents.
36. A method for searching for product offers, comprising:
- 15 providing a computer system adapted to operation of agents;
- providing a consumer agent, each associated with and capable of communicating with a consumer;
- providing a consumer data base of consumer preference data associated with each said consumer;
- 20 providing a data base of offers to sell and offers to buy;
- concealing the identity of said consumer from agents not associated with said consumer;
- said consumer agent assisting said consumer in composing a query for a product or product category;
- 25 said consumer agent searching said data base of offers to sell and offers to buy, wherein said searching generates persistent market data;

-92-

said consumer agent retaining search results of said searching until the
associated consumer's preferred delivery time;
said consumer agent automatically filtering out any of said search results
that violate any of said consumer preference data;
5 said consumer agent ordering or ranking said search results according to
said consumer preference data;
said consumer agent formatting said search results for delivery to said
associated consumer's preferred communication device; and
said consumer agent delivering said search results to said communication
10 device;
whereby said consumer can gather product data that is automatically
filtered and ranked according to said consumer preferences.

37. The method of claim 36 further comprising:

15 providing a data base of advocate evaluation and recommendation data;
said consumer selecting data from said data base of advocate evaluation and
recommendation data to include in said consumer's said consumer
data base of consumer preference data;
said consumer agent thereby further automatically filtering out any of said
20 results that violate consumer's selected advocate evaluation and
recommendation data;
whereby said consumer has a convenient way to use advocate evaluation
and recommendation data to automatically filter and rank said
search results.

25

38. A method for collecting and analyzing market data, comprising:
providing a computer system adapted to operation of agents;

-93-

- providing a plurality of consumer agents, each associated with and capable
of communicating with a consumer;
- providing a consumer data base of consumer preference data associated
with each said consumer;
- 5 providing a plurality of provider agents, each associated with and capable
of communicating with a provider;
- providing a data base of offers to sell and offers to buy;
- concealing the identity of said consumer from agents not associated with
said consumer;
- 10 searching by said consumer agents of said data base of offers to sell and
offers to buy, wherein said searching generates persistent market
data;
- searching by said provider agents of said consumer data base and of said
persistent market data to discover which of said consumer agents
- 15 possess certain preferences and generated certain records within said
persistent market data;
- analyzing the results of said searching by said provider agents to quantify
consumer demand;
- whereby consumers are encouraged to utilize said system because the
- 20 identity of each said consumer is concealed from said agents not
associated with said consumer in said system; and
- whereby providers can gain access to valuable said persistent market data
without imposing on or invading the privacy of said consumers.
- 25 39. The method of claim 38 wherein said searching by said provider agents is
restricted to current data within said persistent market data, and said
analyzing quantifies current consumer demand.

-94-

40. The method of claim 38 wherein said searching by said provider agents is restricted to historical data within said persistent market data, and said analyzing quantifies historical consumer demand.
- 5 41. A method for quantifying future demand for a certain product or a certain product category, utilizing persistent market data generated by consumer agents that conceal the identity of said consumer agent's associated consumer, while searching a data base of offers to sell and offers to buy, and utilizing a consumer data base of consumer preference data,
- 10 comprising:
searching said consumer data base to discover which of said consumer agents possess certain preferences;
searching said persistent market data to discover which of said consumer agents, while searching for products related to said certain product or certain product category, generated certain records within said
- 15 persistent market data;
whereby a provider can select specific consumers that may be interested in said certain product or certain product category.
- 20 42. A method for targeting specific consumers, each of whose identity remains concealed, according to their on-line shopping activities and preferences, utilizing persistent market data generated by consumer agents that conceal said identity while searching a data base of offers to sell and offers to buy, and utilizing a consumer data base of consumer preference data,
- 25 comprising:
searching said consumer data base to discover which of said consumer agents possess certain preferences;

-95-

searching said persistent market data to discover which of said consumer agents generated certain records within said persistent market data; whereby a provider can select specific consumers to receive messages.

- 5 43. The method of claim 42 wherein said certain records pertain to said consumer agents searching for specific products.
44. The method of claim 42 wherein said certain records pertain to said consumer agents searching within a product category.
- 10 45. A method of extending a consideration payment to a consumer, whose identity remains concealed, in return for said consumer providing useful market data to a provider, comprising:
- said provider selecting consumer agents to receive a message;
- 15 said provider causing said message to be delivered to said consumer agents; each of said consumer agents displaying said message to said consumer agent's associated consumer;
- each of said associated consumers generating a response to said message; said provider's associated provider agent generating, for each said response
- 20 received, a consideration notice addressed to the consumer agent associated with the consumer that generated the response;
- said associated provider agent causing each said consideration notice to be delivered to said associated consumer agent;
- each said consumer agent crediting said associated consumer's
- 25 consideration account;
- whereby said consumer can receive a consideration payment for divulging useful market data to said provider; and

-96-

whereby said provider can solicit valuable market data without imposing on
or invading the privacy of said consumers.

46. The method of claim 45 wherein said message is an advertisement.

5

47. The method of claim 45 wherein said message is a marketing survey.

48. The method of claim 45 wherein said message is a questionnaire about the
reason for a sale.

10

49. The method of claim 45 wherein said message is a questionnaire about the
reason for a lost sale.

50. A method of a consumer automatically rejecting certain unsolicited
messages from a provider, comprising:
said provider causing a message to be delivered to a consumer agent;
said consumer agent comparing aspects of said message to a consumer data
base of consumer preference data;

15

if said message satisfies said consumer preference data, said consumer
agent allowing said message to complete delivery to said consumer;
if said message violates said consumer preference data, said consumer
agent rejecting said message by:

20

automatically generating a rejection message including an indication
of the violated consumer preferences;

25

causing said rejection message to be delivered to said provider;
whereby said provider can gain valuable market data even from messages
that are blocked from delivery.

-97-

51. A method of a provider simulating demand, comprising:
providing a data base of offers to sell and offers to buy;
providing a provider agent associated with and capable of communicating
with a provider;
5 providing a plurality of consumer agents, each associated with and capable
of communicating with a consumer;
said provider places an invisible advertisement in said data base of offers to
sell and offers to buy;
a consumer agent, when ordering or ranking search results, if said search
10 results include a reference to said invisible advertisement, ranks said
invisible advertisement along with other said search results;
said consumer agent, when ranking said invisible advertisement, generates
a message indicating the ranking of the invisible advertisement;
said consumer agent sends said message to the associated provider agent;
15 said consumer agent, when delivering said search results to the preferred
communication device of the associated consumer, omits said
invisible ad;
whereby said provider can determine simulated demand for the product
described by said invisible ad.
20
52. A method of a provider replaying demand, comprising:
providing a plurality of consumer agents, each associated with and capable
of communicating with a consumer;
providing a data base of offers to sell and offers to buy;
25 consumer agents searching said data base of offers to sell and offers to buy,
whereby said consumer agent searching generates persistent market
data;

-98-

a provider composes an invisible advertisement;
said provider selects at least one said consumer agent that has completed
said searching;
said consumer agent is instructed to rank said invisible advertisement;
5 said consumer agent ranks said invisible advertisement along with other
said search results;
said consumer agent, when ranking said invisible advertisement, generates
a message indicating the ranking of the invisible advertisement;
said consumer agent sends said message to the associated provider agent;
10 whereby said provider can determine simulated demand in historical time
for the product described by said invisible ad.

53. In a computer network agent system for providing communication between
an anonymous potential consumer of products which can be goods or
15 services and a provider of such products, the combination comprising:
a decision agent for receiving anonymous product queries from the
consumer and transmitting product recommendations to the
consumer;
a demand agent for receiving demand queries from the provider and
20 transmitting quantified demand information to the provider; and
a market for gathering information from the agents, organizing the
information and distributing organized information to the agents.

54. A combination according to Claim 53 wherein the decision agent
25 comprises, in combination:
a unique identifier function for maintaining an identifier that uniquely
identifies this decision agent within the agent system;

-99-

- a market reference for indicating in which market the decision agent should search;
- an expiry function for indicating how long the decision agent should continue searching;
- 5 a query for describing the product or product category for which to search;
- a response manager for receiving search results and returning the search results to the consumer; and
- a log function for storing records of the activities of the decision agent.
- 10 55. A combination according to Claim 53 wherein the demand agent comprises, in combination:
- a unique identifier function for maintaining an identifier that uniquely identifies this demand agent within the agent system;
- a market reference for indicating in which market or markets the demand agent should search;
- 15 a datetime range function for indicating that demand should be quantified over the date/time range specified;
- a demand query for describing a product or product category query that can be matched against the queries of decision agents; and
- 20 a log function for storing records of the activities of the demand agent for later consultation by other components of the agent system.
56. A combination according to Claim 53 wherein the market comprises, in combination:
- 25 a product listing function for maintaining a list of the products that can be advertised in this market;

-100-

- a cross reference manager for maintaining cross references to other markets that carry similar products;
- a sell advertisement manager for accepting advertisements of offers to sell that are submitted by consumers;
- 5 a buy advertisement manager for accepting advertisements of offers to buy that are submitted by consumers;
- an active demand agent manager for maintaining a list of all demand agents that are currently calculating demand in this market;
- a template dispenser for retrieving data that is available within the agent system about a particular product;
- 10 a query logger for archiving summary information about queries so that historical data about queries may be quickly accessed without having to access the detailed data which has been archived; and
- an historical demand search engine for matching demand queries of demand agents against the queries that have been previously logged by the query logger to identify decision agents that previously, during a specified datetime range, had searched for a certain product or product category.
- 15
- 20 57. A combination according to Claim 56 wherein the market further comprises a remote database adaptor for providing communication and session management services to connect to a remote database belonging to a manufacturer or a provider and translating between the data formats used by the remote database and the data formats used by the market.
- 25

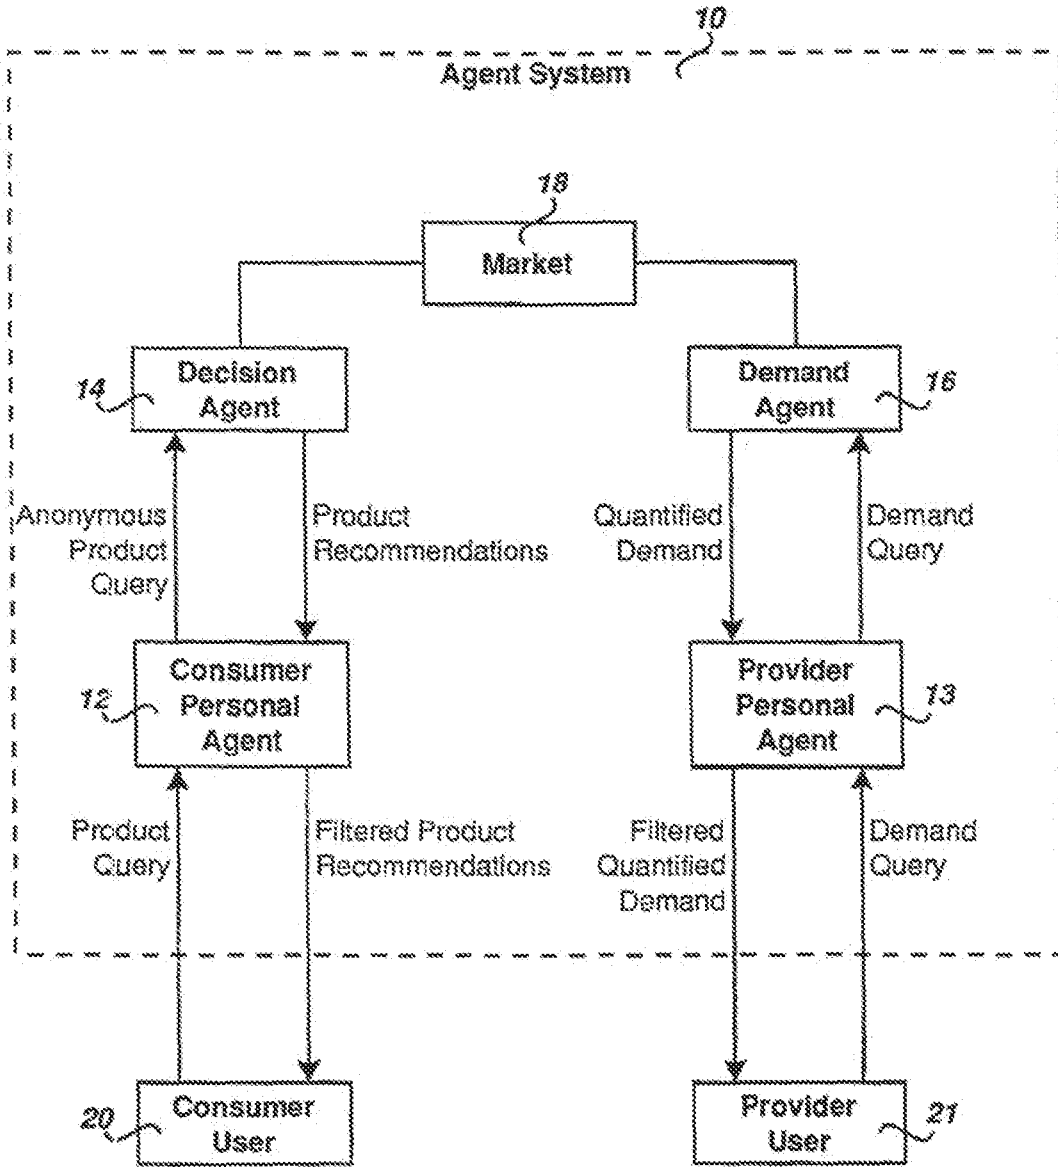


Fig. 1

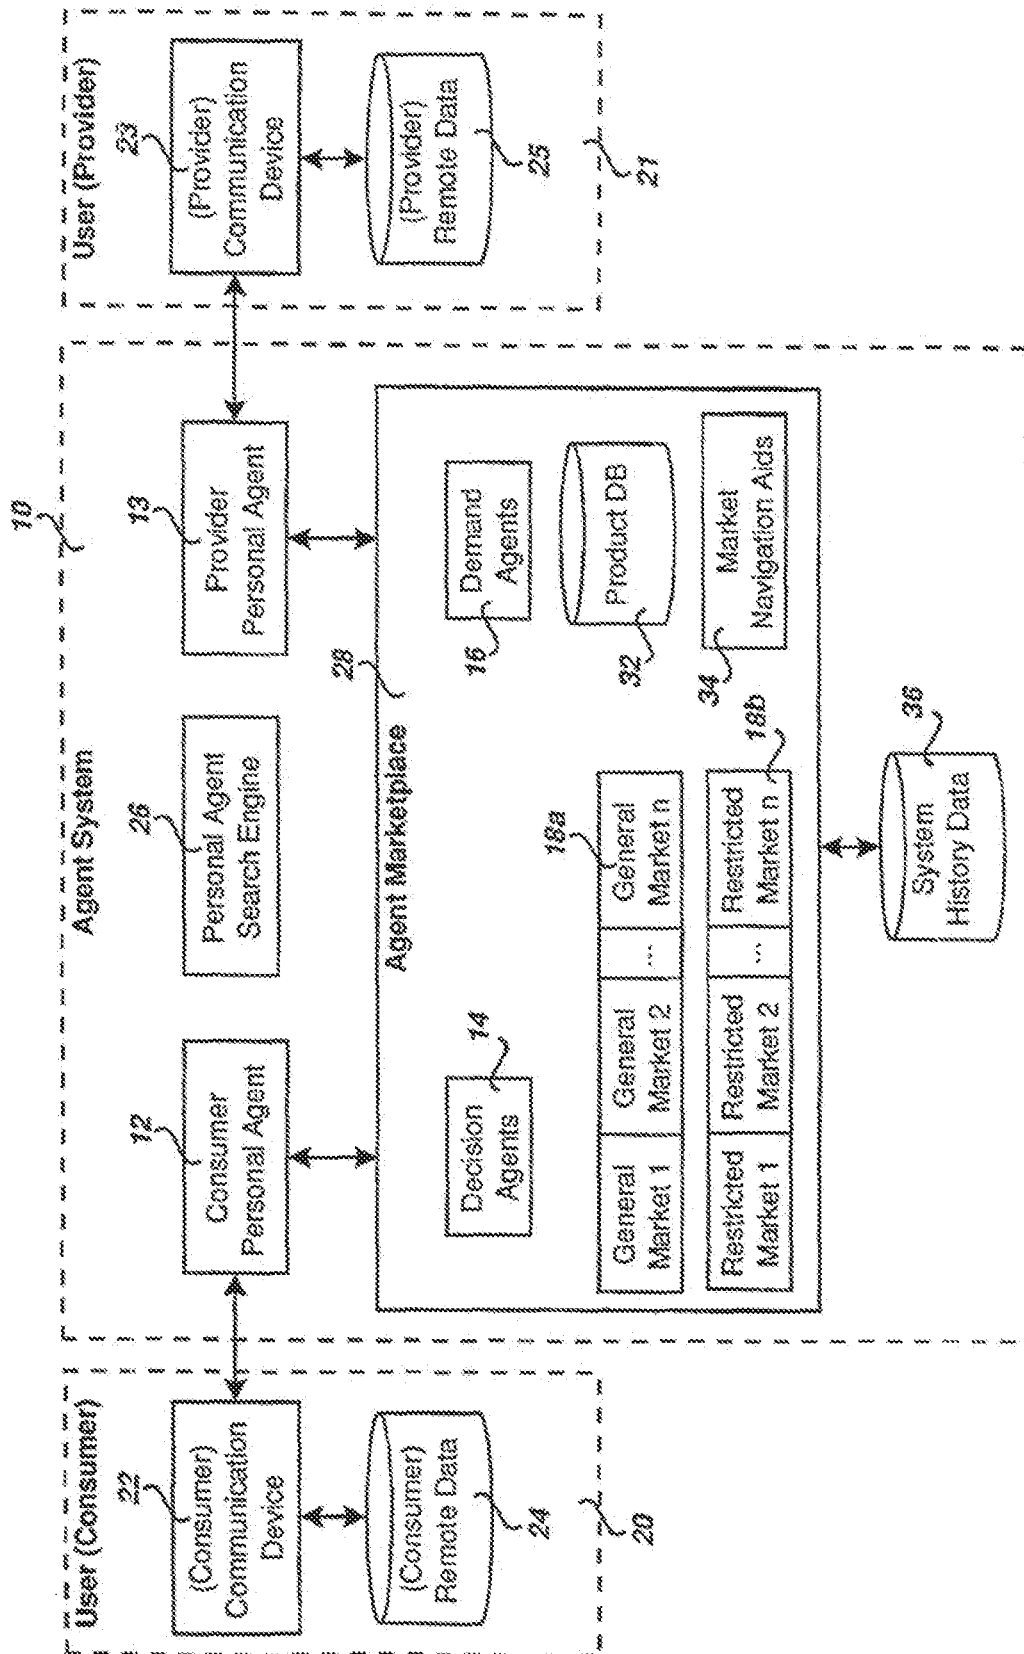


Fig. 2

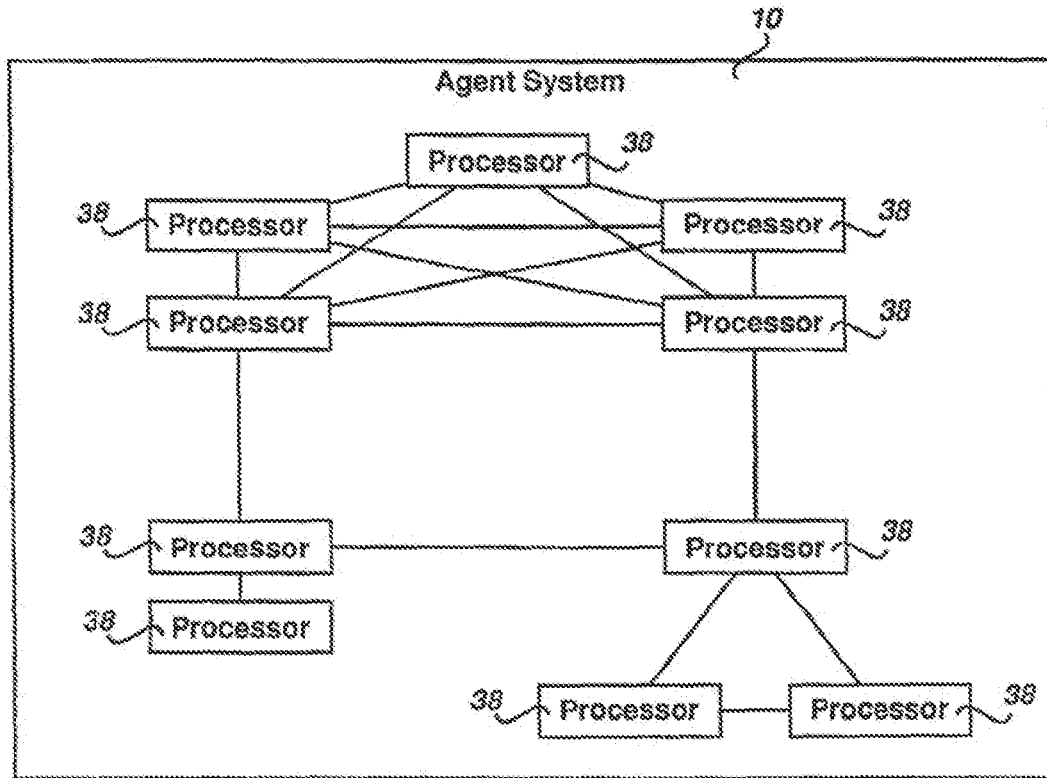



Fig. 3A

Fig. 3B






Fig. 4A




Fig. 4B




Fig. 4C




Fig. 4D




Fig. 5A

Key	Value
Age	34
Homeowner	Yes
Gender	Male
Cats	interested
brand name 1	like
brand name 2	dislike
brand name 3	neutral
brand name 4	like > brand name 3
brand name 5	a favorite
email Consideration Fee	greater than \$1.00
alpine skiing	dislike
cross country skiing	like
MSG in food	dislike
delivered pizza	No
phone solicitation	never
favorite color	blue, red
health and fitness	interested
weight lifting	rank 1 in 10
stair climbing	rank 3 in 10
swimming	rank 10 in 10

Fig. 5B




Fig. 6





Fig. 7

Fig. 8A



118




Fig. 8B

152




Fig. 8C

Fig. 9A

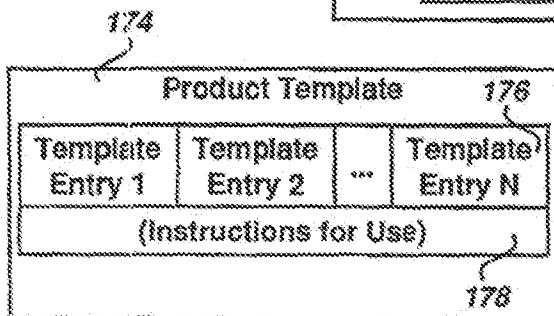
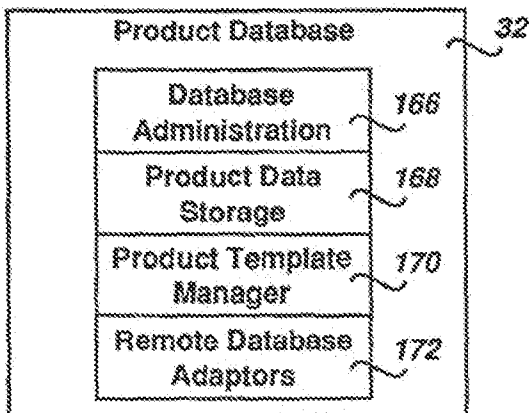


Fig. 9B




Fig. 9C

<i>Keyword</i>	<i>Specification</i>	<i>Use Flags</i>	<i>Value</i>
Product	enumeration	search,ad	Television Set
Brief Description	text	ad	27" TV
Brand	enumeration	search,ad	<i>brand name</i>
Model	text	search,ad	392-43
Mfr Sug Retail Price	money	search,ad	\$500.00
Screen Size	number	search,ad	27
Remote Control	boolean	search,ad	yes
Cable Ready	boolean	search,ad	yes
Stereo Sound	boolean	search,ad	yes
Cabinet Color	enumeration	search,ad	black
Expanded Info	text	ad	Futuristic design...
Image	image	ad	(picture of product)
<i>rating guide</i> Rating	rank	search	3 of 14
Endorser	text	search	<i>endorser</i>

Fig. 9D





Fig. 10

11/46

Product Search Method

Fig. 11



12/46

Compose Decision Query Subroutine

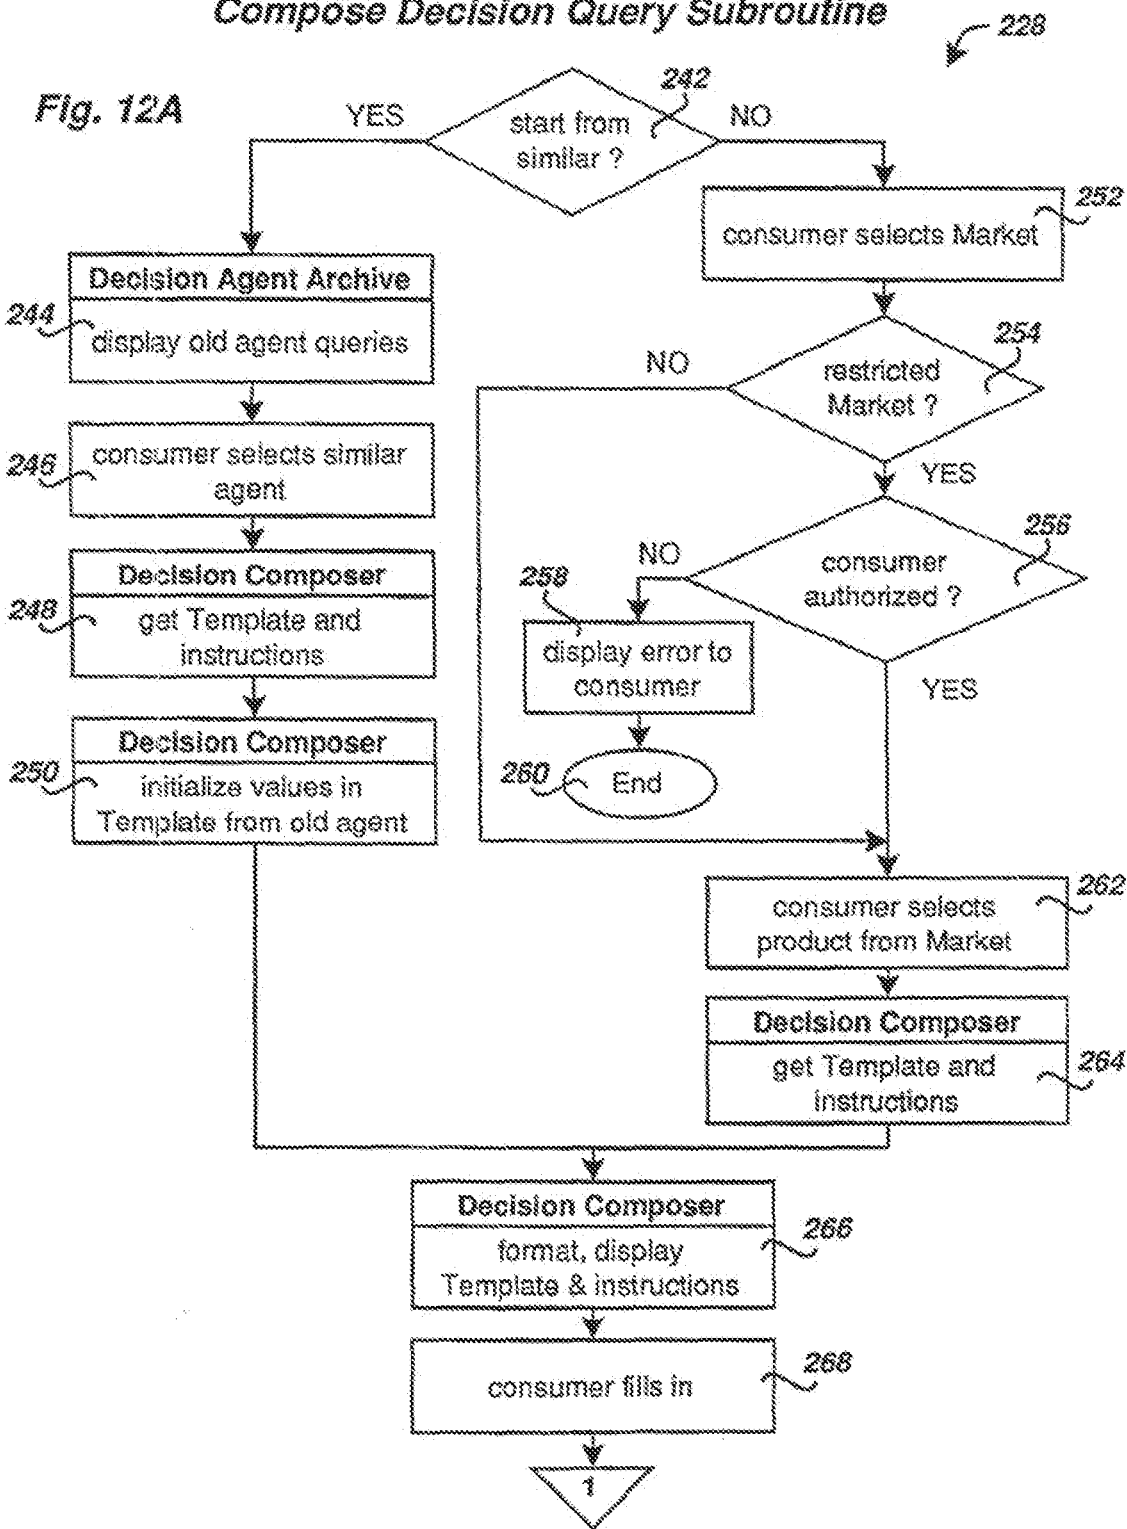

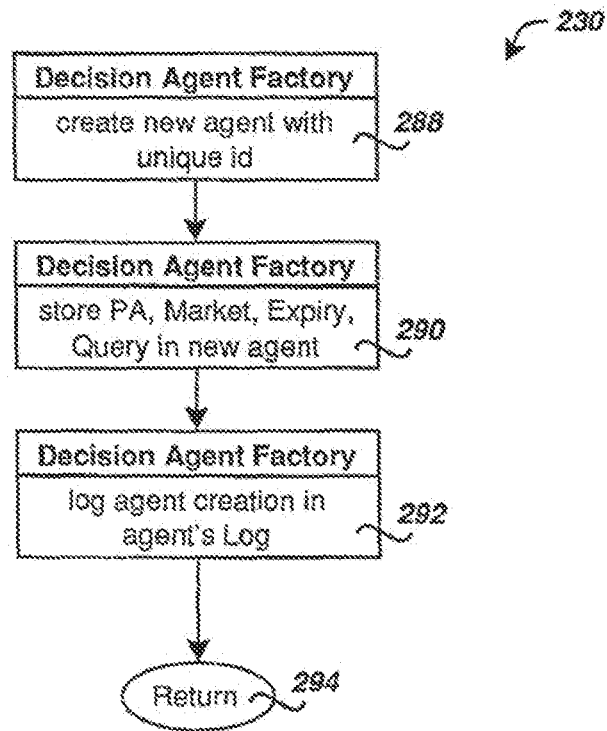


Fig. 12B




Create Decision Agent Subroutine

Fig. 13



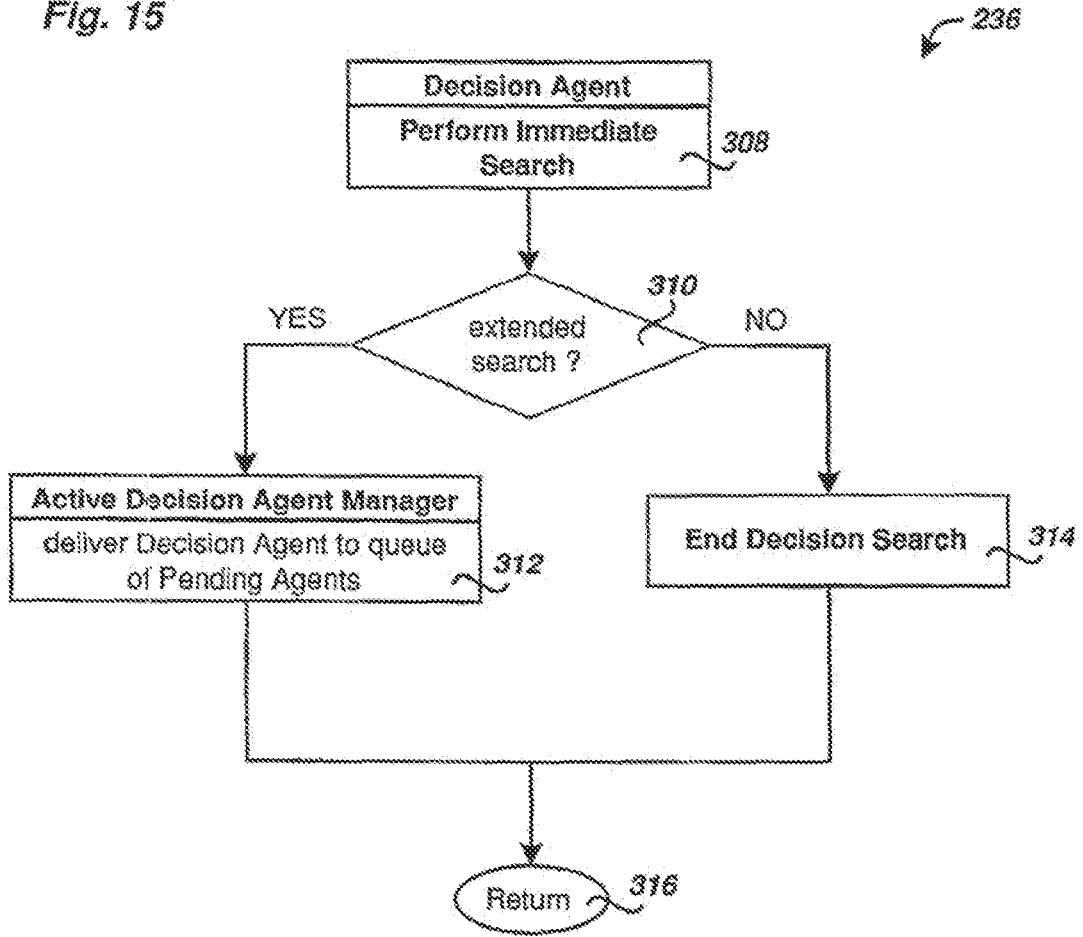
Accept New Decision Agent Subroutine

Fig. 14




Perform Decision Search Subroutine

Fig. 15



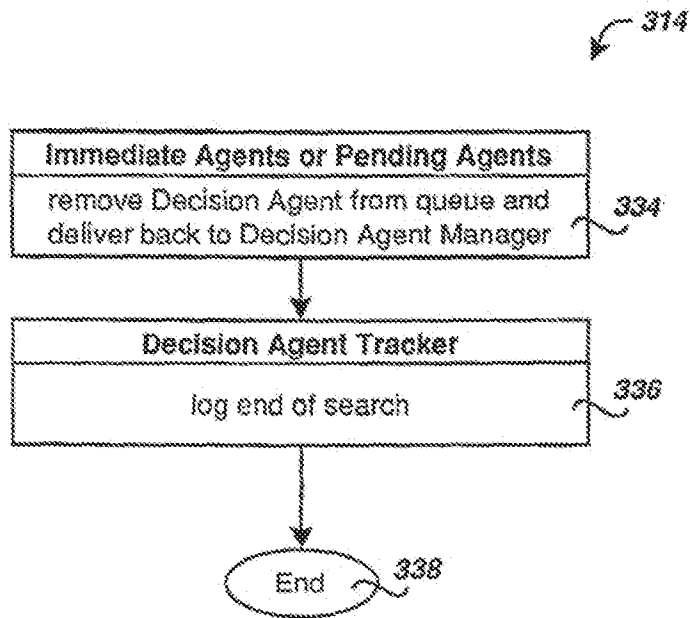
Perform Immediate Search Subroutine

Fig. 16



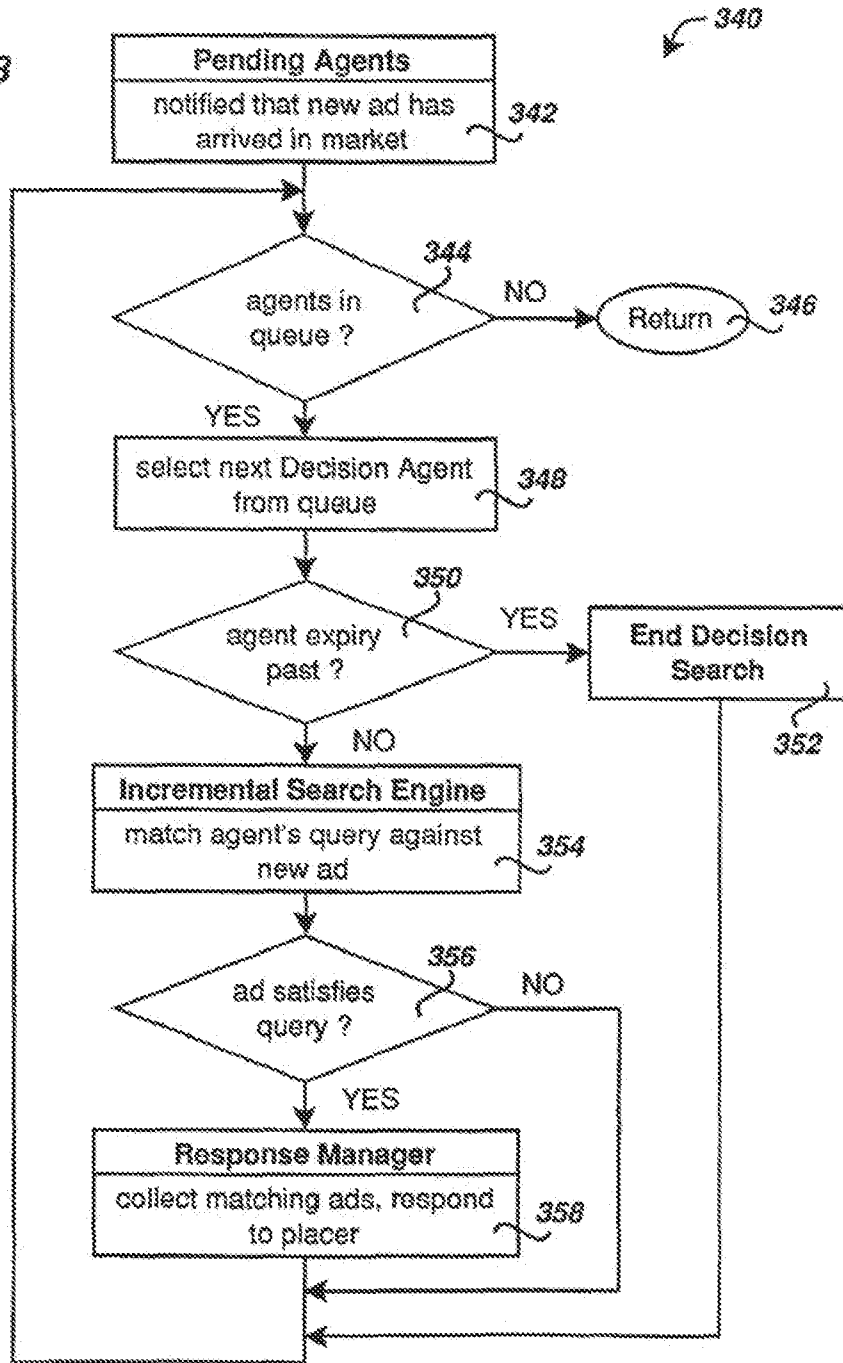
End Decision Search Subroutine

Fig. 17



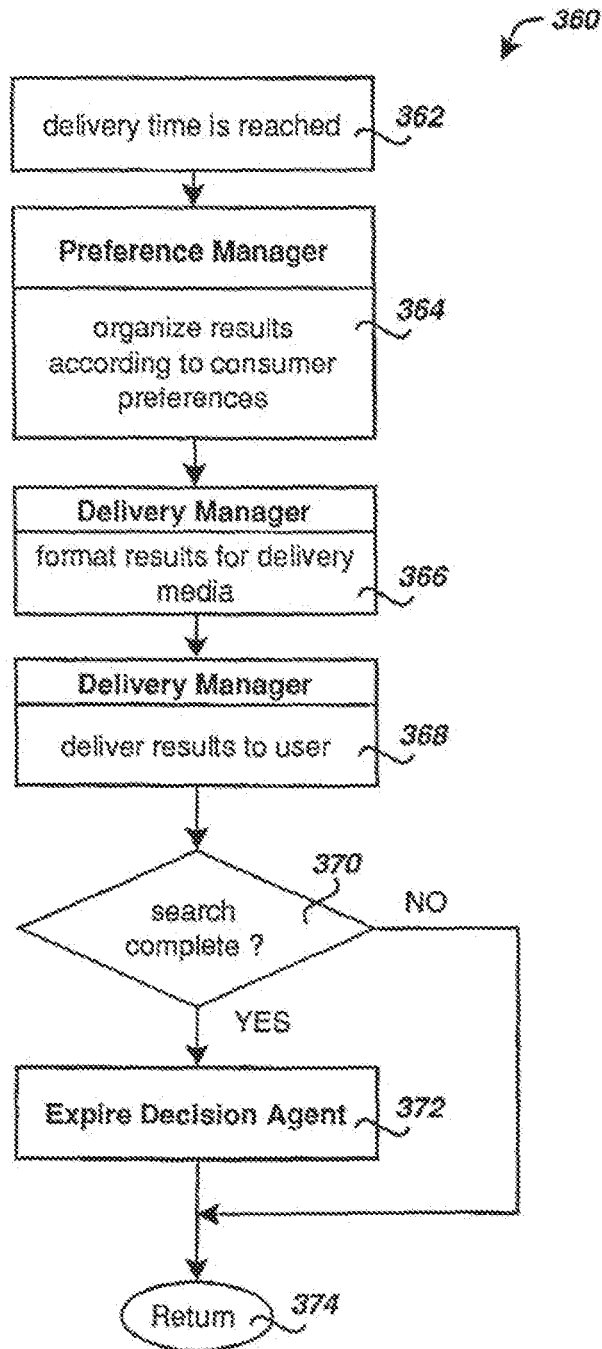
Extended Search Subroutine

Fig. 18



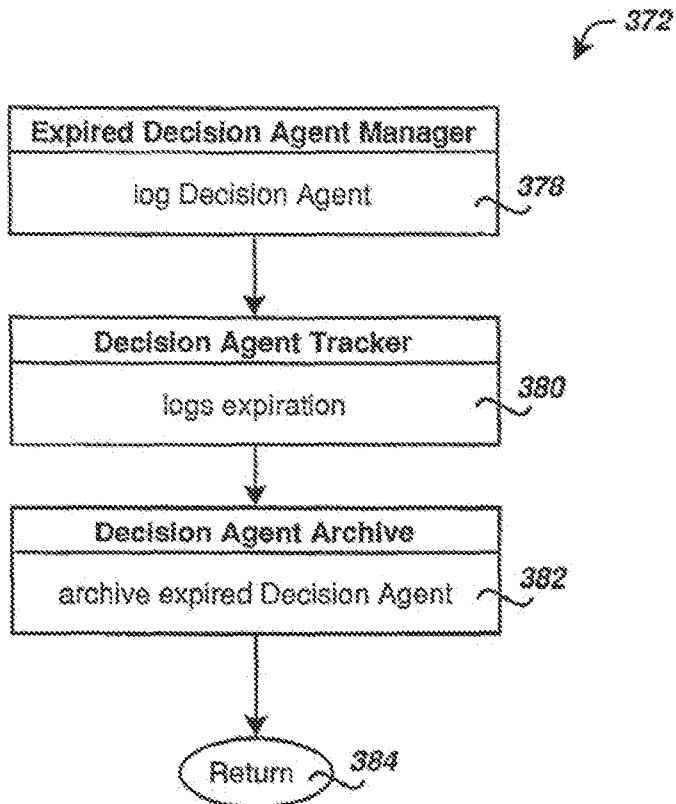
Deliver Search Results Subroutine

Fig. 19




Expire Decision Agent Subroutine

Fig. 20



Quantify Demand Method

Fig. 21



Compose Demand Query Subroutine

Fig. 22A

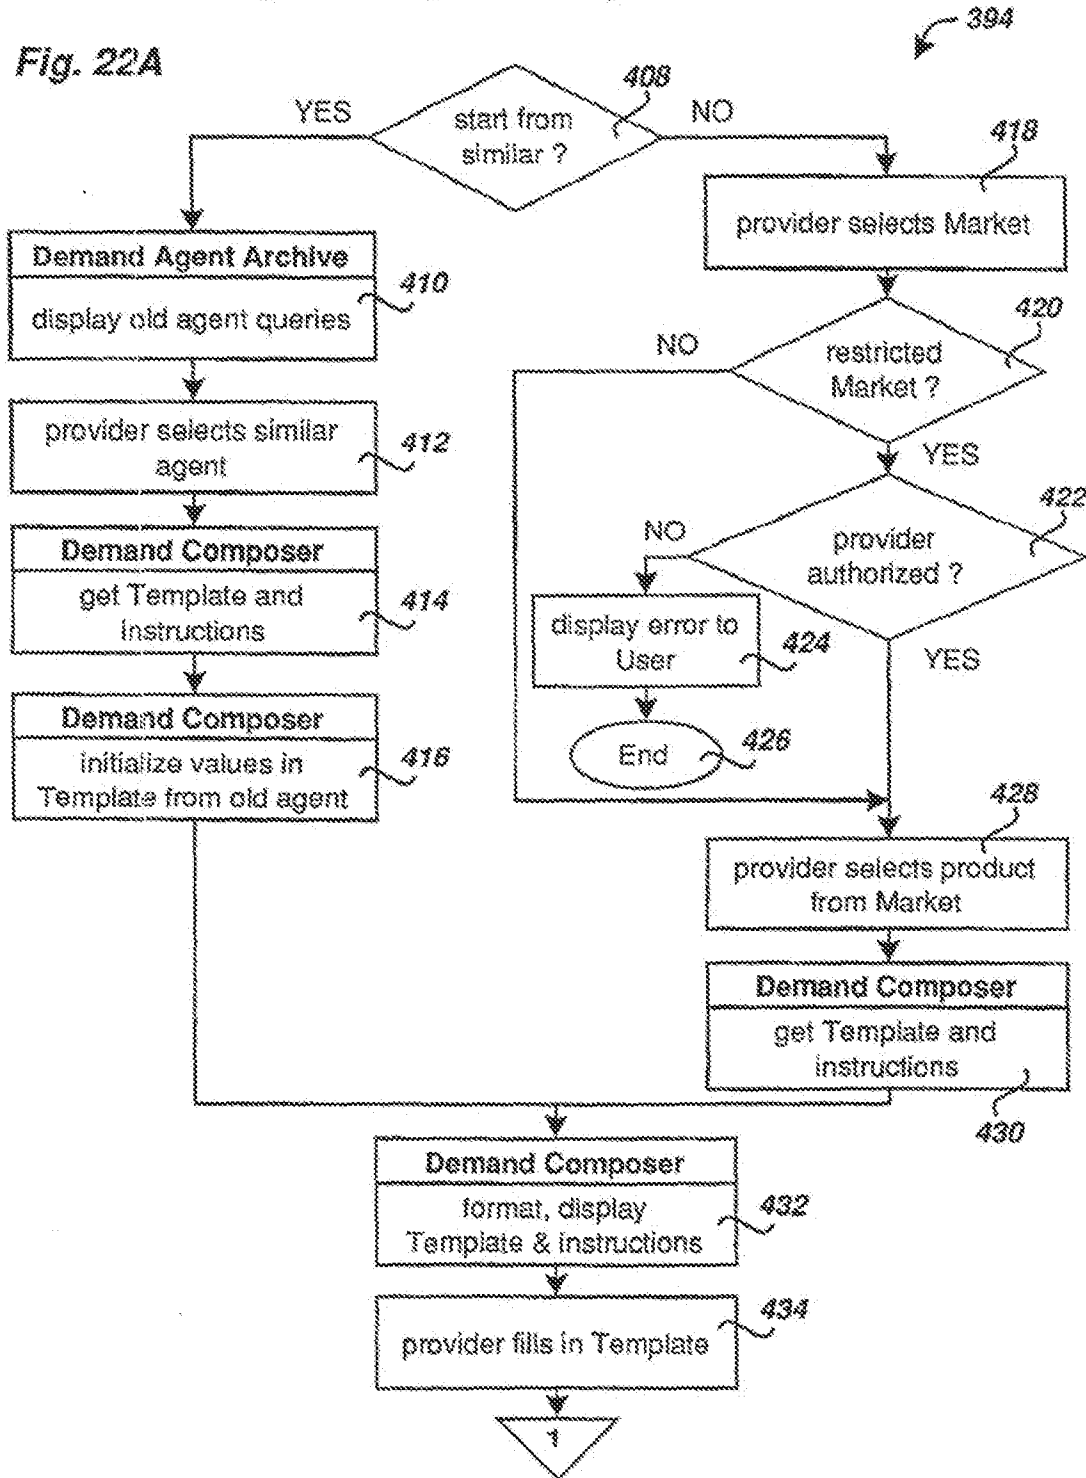
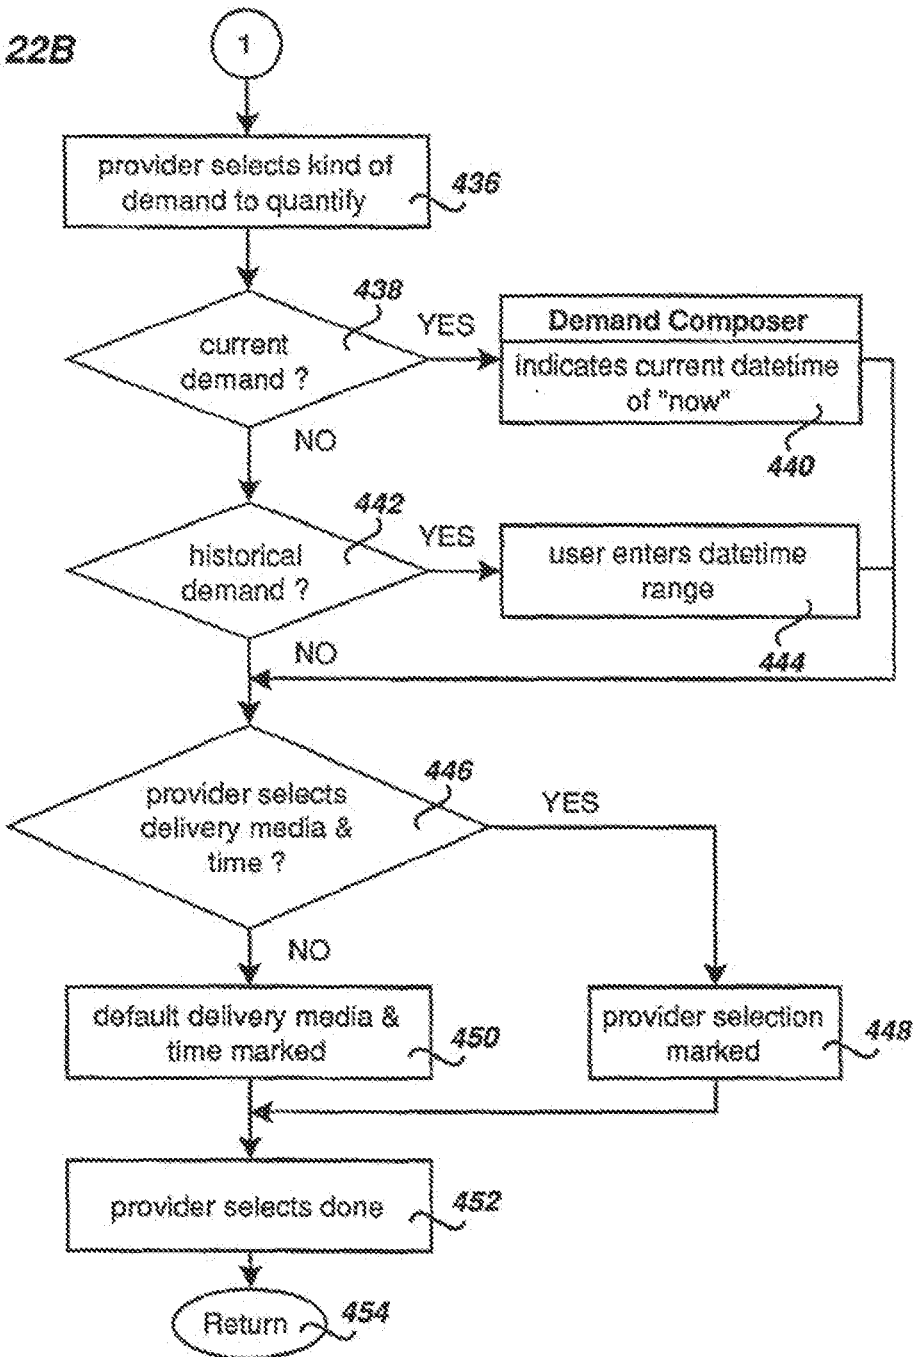
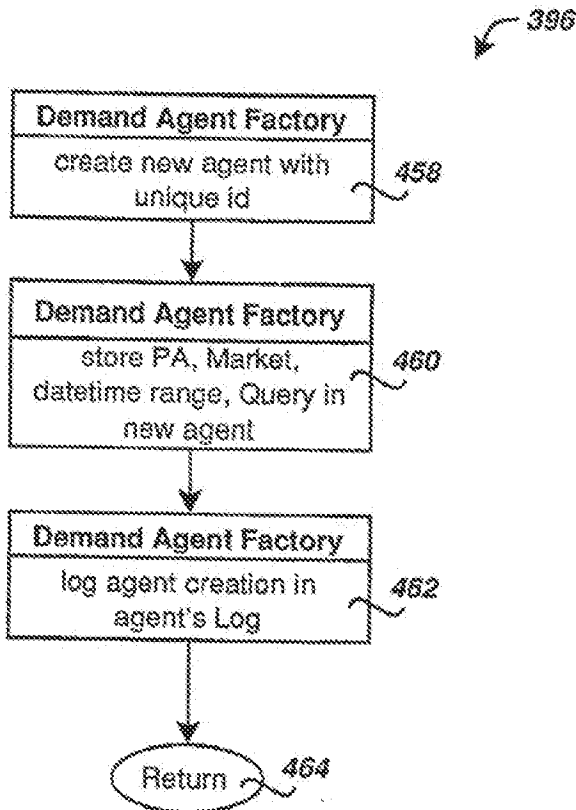


Fig. 22B



Create Demand Agent Subroutine

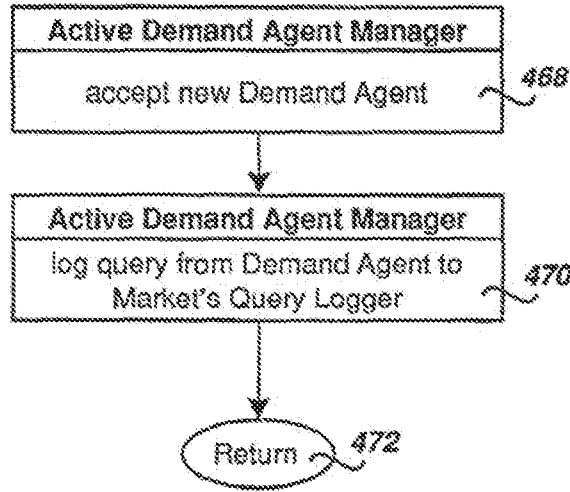
Fig. 23



Accept New Demand Agent Subroutine

Fig. 24

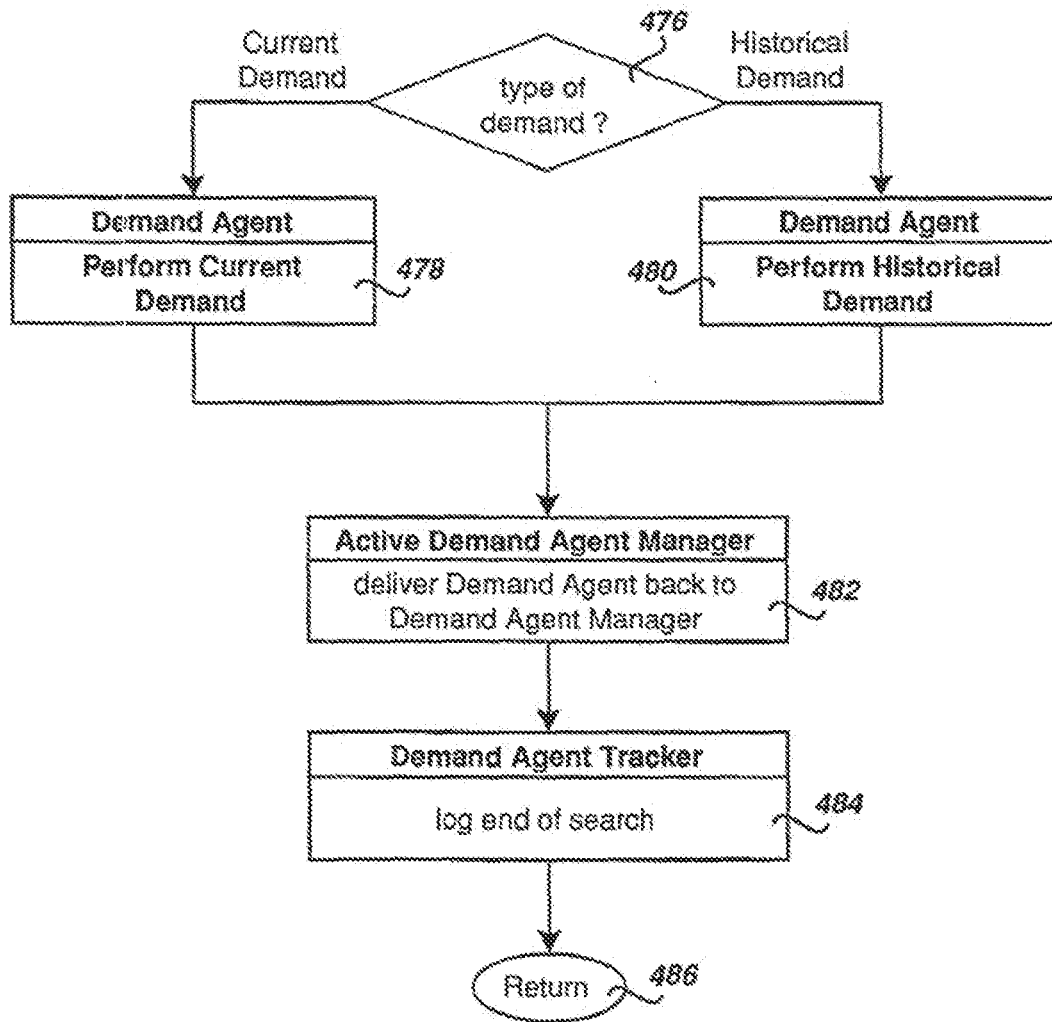
400



Perform Demand Search Subroutine

Fig. 25

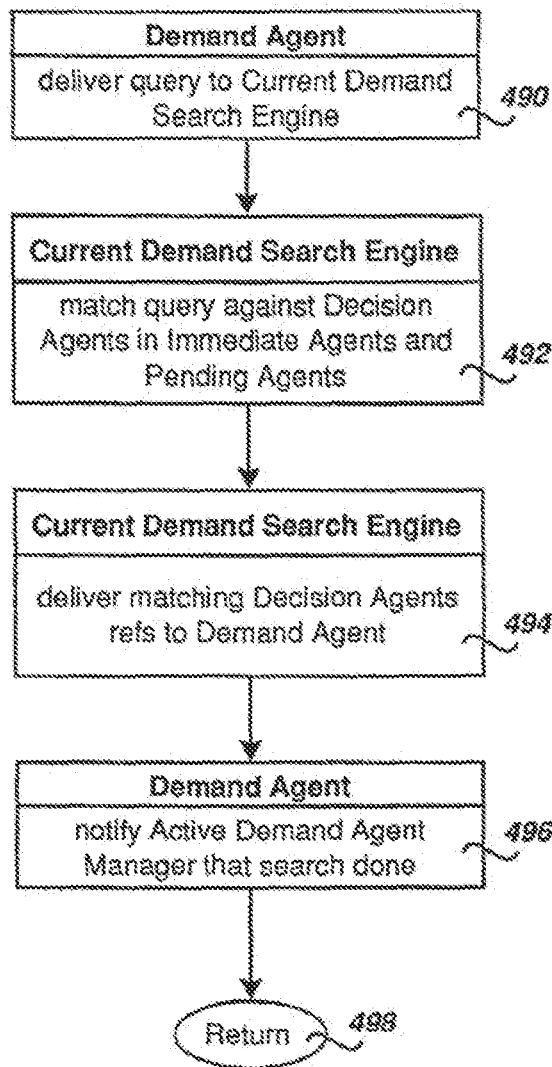
402



Perform Current Demand Subroutine


Fig. 26

478



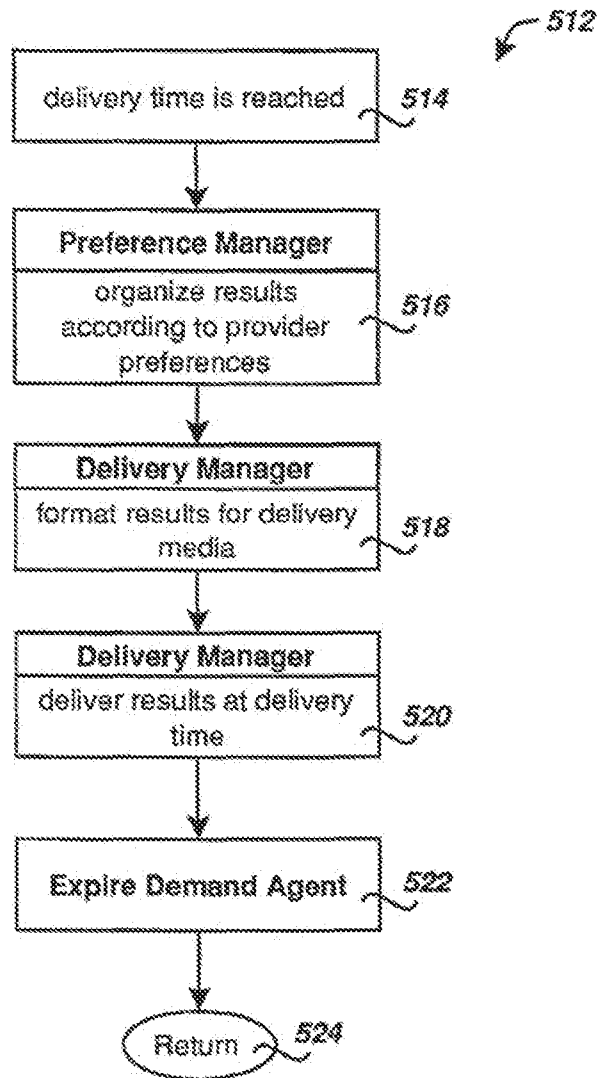
Perform Historical Demand Subroutine

Fig. 27



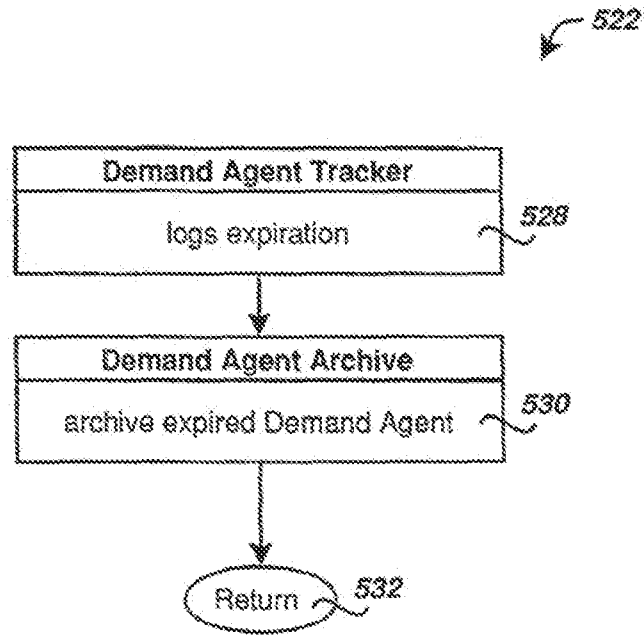
Deliver Demand Results Subroutine

Fig. 28



Expire Demand Agent Subroutine

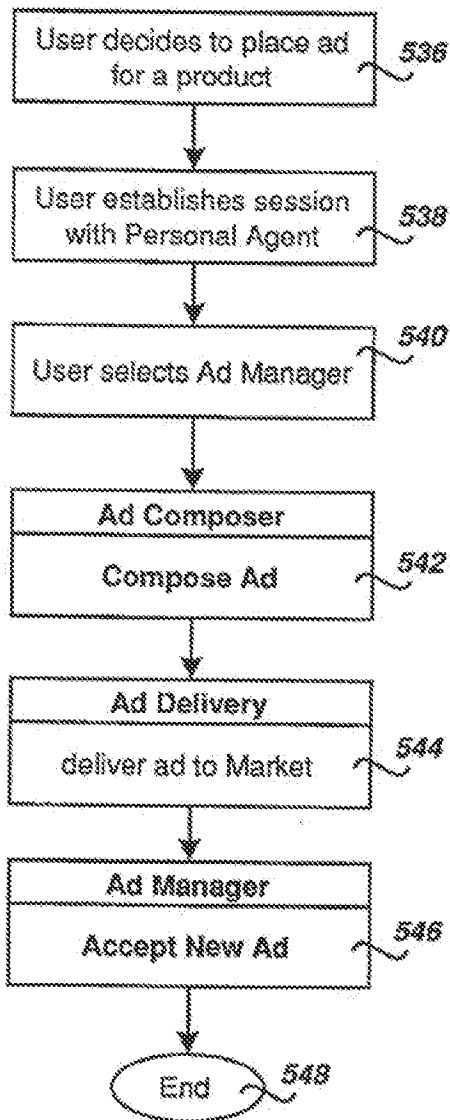
Fig. 29



Place Ad Method

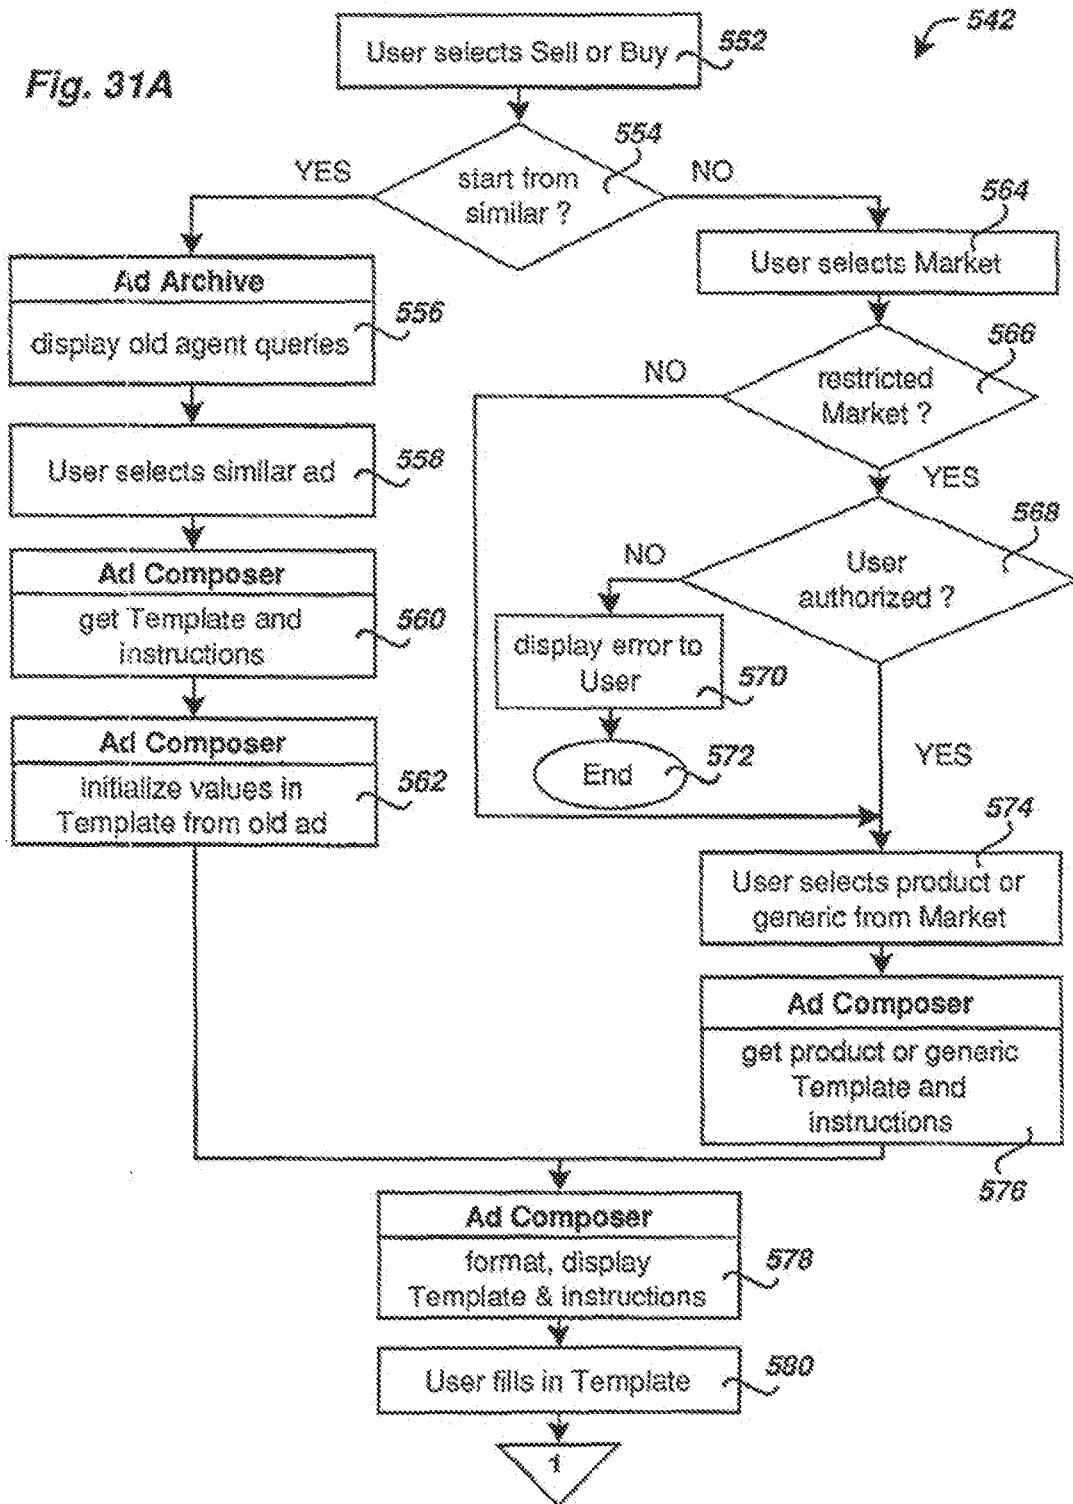
Fig. 30

534



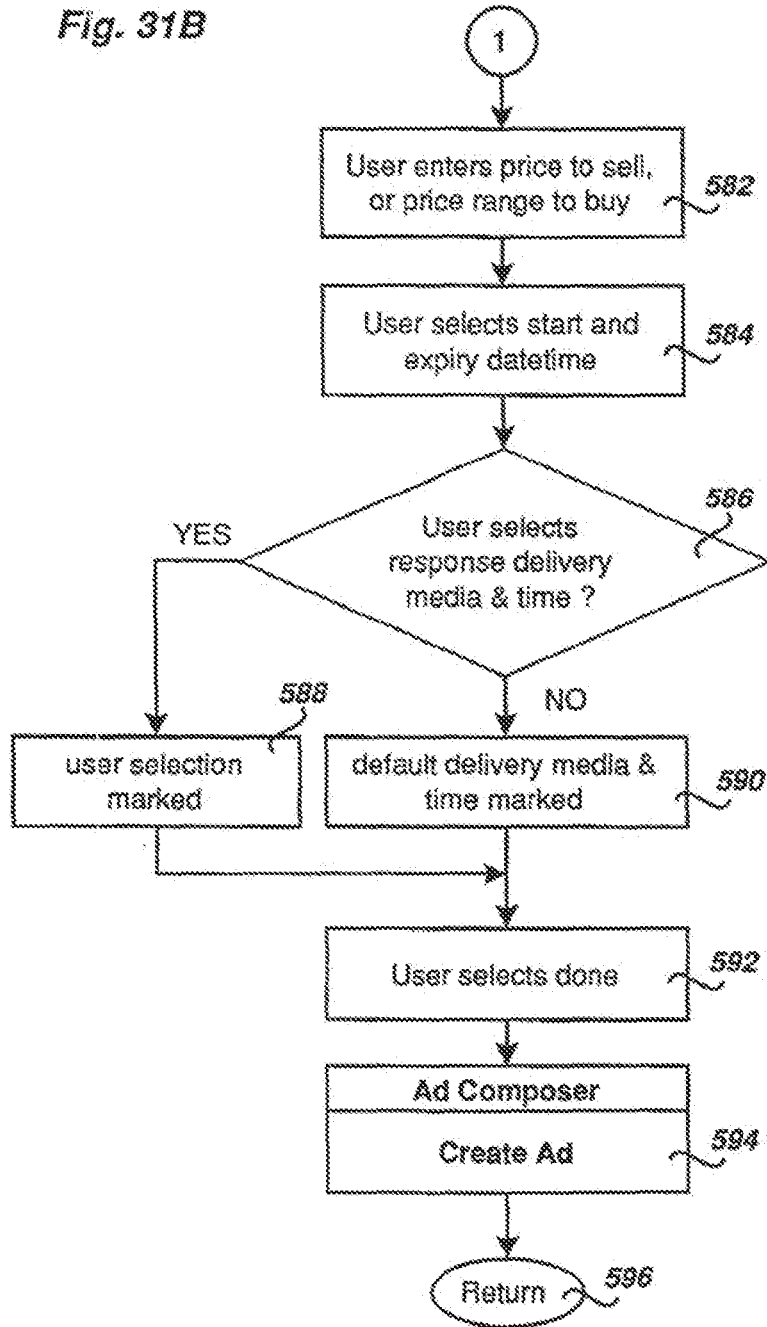
Compose Ad Subroutine

Fig. 31A



34/46


Fig. 31B



Create Ad Subroutine


Fig. 32

594




Accept New Ad Subroutine

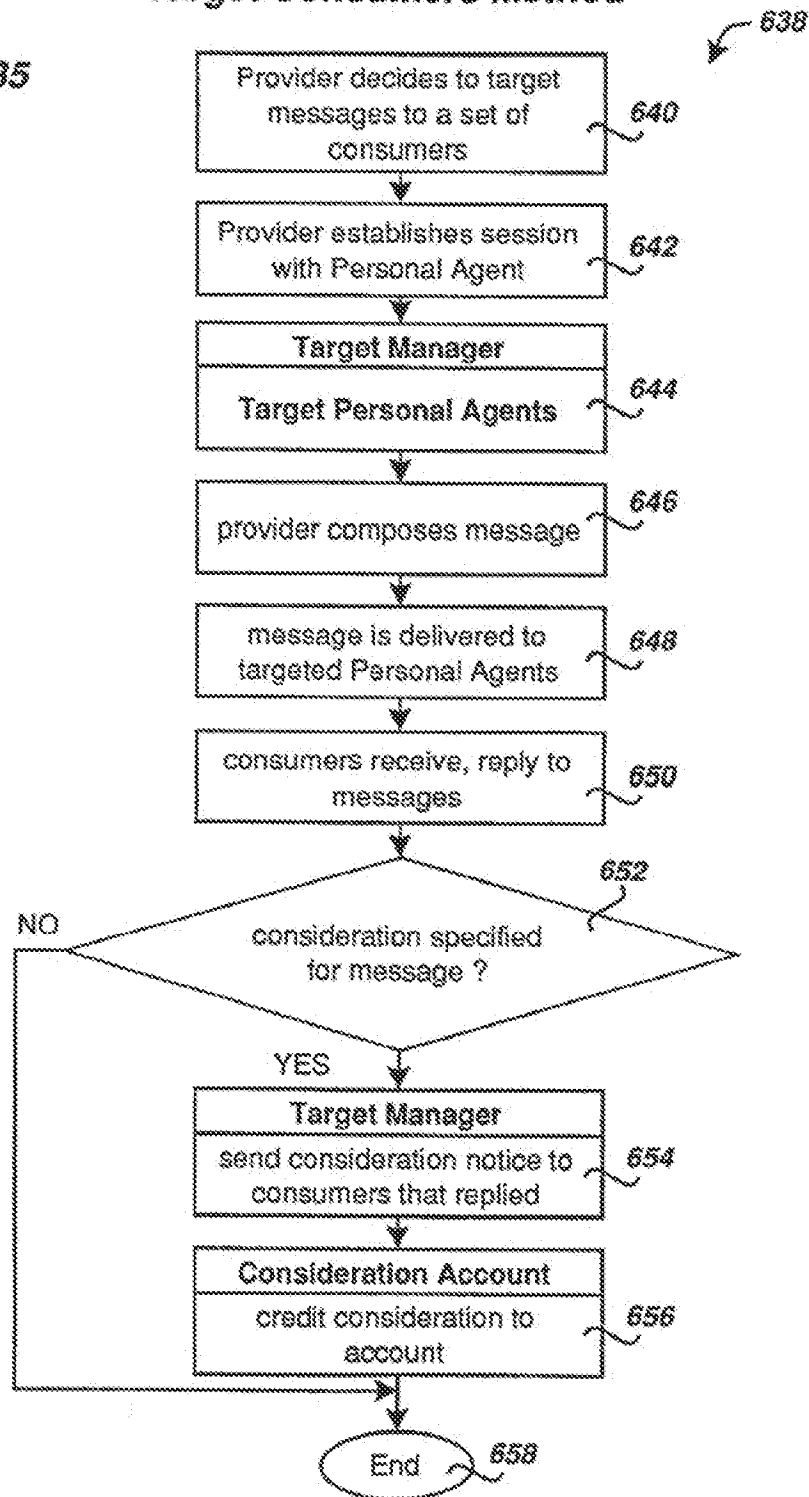
Fig. 33



Expire Ad Subroutine

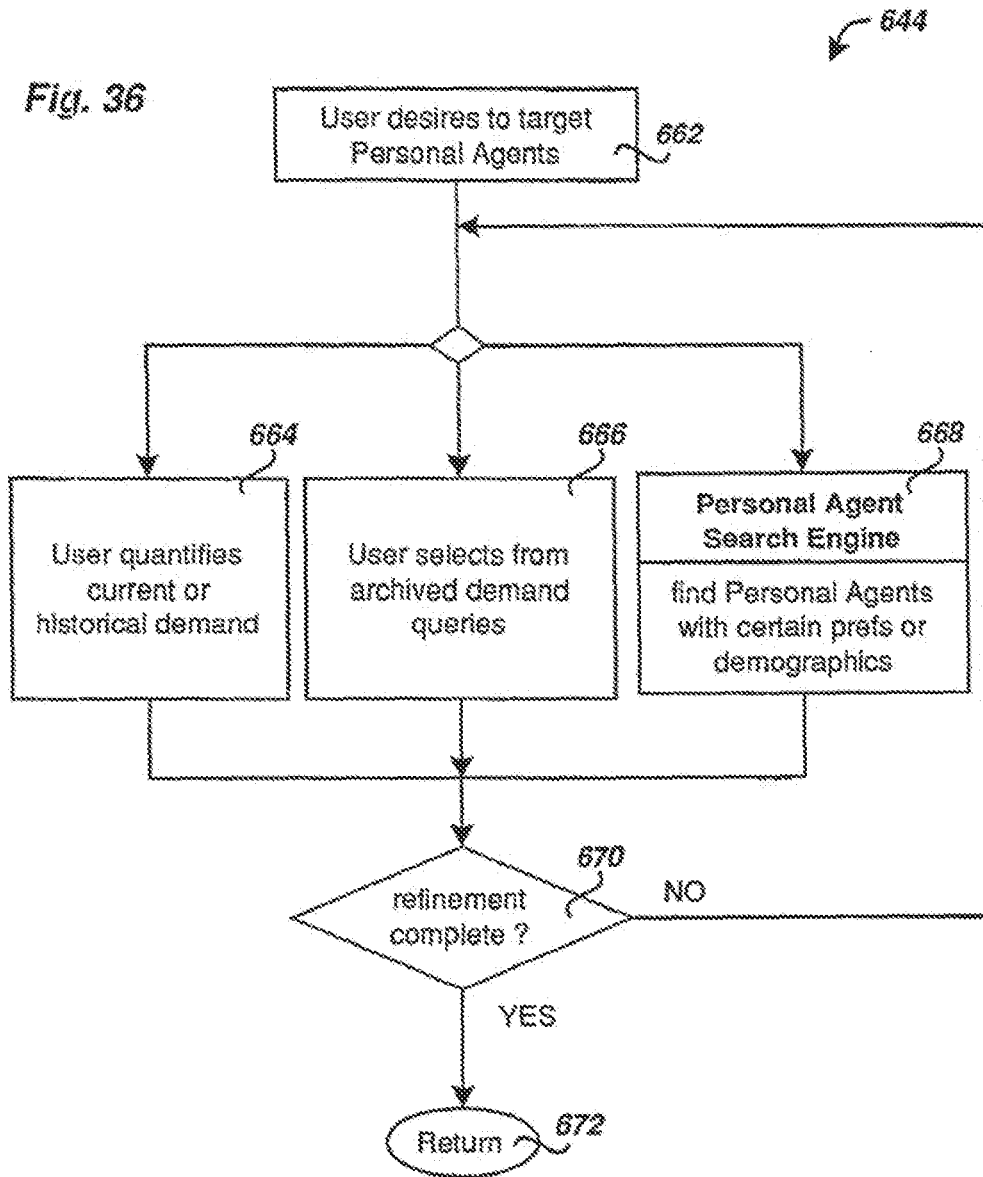
Fig. 34



Target Consumers Method**Fig. 35**


Target Personal Agents Subroutine

Fig. 36




Reject Unsolicited Message Method

Fig. 37

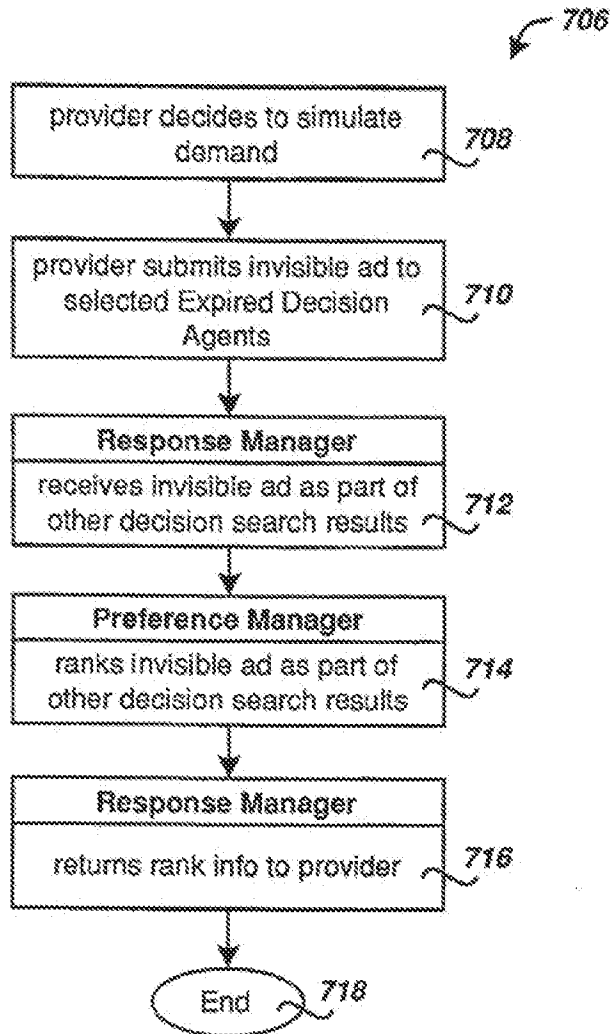


Simulate Demand Method

Fig. 38A



42/46

Replay Demand Method**Fig. 38B**

Login to Agent System

Agent Logo *Agent System Name*

Welcome to Agent System

What is a Personal Agent? Create your own for FREE!

If you already have a Personal Agent,

enter your Agent Id:

and Password:

[then Click here to enter the system](#)

Fig. 39

Search the Consumer Electronics Market

Need detailed instructions? [Click here](#)

icon

Search for Consumer Electronics

Tell us what you're looking for, and let your Personal Agent immediately search for you!

Category

TV

VCR

Laser Disk Player

Cassette Player

Cassette Recorder

Compact Disc Player

Complete Stereo System

Speakers

Receiver

Amplifier

Tuner

Game Systems

Clock Radio

Radio

Accessories

Component

Portable

Console

Product	Brand:		Model:	
Location	Merchant:		City:	
Price Range	From:		Up to:	

Click here to specify additional features

Click here to start looking

Fig. 40

Search the Automobile Market

Need detailed instructions? [Click here](#)

icon

Search for Automobiles

Tell us what you're looking for, and let your Personal Agent immediately search for you!

Sedan

Mini Van

Full size Van

Sport Utility

Mini Pickup

Full size Pickup

High Performance

Luxury

Convertible

2 door

4 door

Hatch back

4 wheel drive

Front wheel drive

Rear wheel drive

Compact

Midsize

Fullsize

New

Used

Vehicle	Make: <i>select</i> <input type="text"/>	Model: <input type="text"/>
Location	Dealer: <input type="text"/>	City: <input type="text"/>
Age	From Year: <input type="text"/>	Thru Year: <input type="text"/>
Price Range	From: <input type="text"/>	Up to: <input type="text"/>
Maximum mileage: <input type="text"/>		

Click here to specify additional features

Click here to start looking

Fig. 41

Compose Ad

Need detailed instructions? [Click here](#)

icon **Compose a Sell Ad for Television Set**

Details from Product Database are automatically included in ad.

Brand name:		Model number:	
UPC code:		<input type="checkbox"/> Identify seller in ad	
Price:		<input type="checkbox"/> Buyer responses to Personal Agent	

If neither box is checked, responses will go to Personal Agent only

Start date & time:	Run through:
--------------------	--------------

Enter additional text for ad. Delivered to searchers who request more info.

[Click here to continue to next step](#)

Fig. 42

(12)

(21) 2 314 513

(51) Int. Cl. 7: G06F 17/60, H04L 12/16, G06F 17/30

(22) 25.07.2000

(30) 09/361,777 US 26.07.1999
09/476,600 US 31.12.1999

(72) BARDY, GUST H. (US).

(71) BARDY, GUST H.,
2518 Constance Dr. W., SEATTLE, XX (US).


(74) OSLER, HOSKIN & HARCOURT LLP

(54) SYSTEME ET METHODE DE NORMALISATION DE LA RETROACTION VOCALE D'UN PATIENT DANS UN SYSTEME AUTOMATISE DE COLLECTE ET D'ANALYSE DE DONNEES SUR LES SOINS AUX PATIENTS
(54) SYSTEM AND METHOD FOR PROVIDING NORMALIZED VOICE FEEDBACK FROM AN INDIVIDUAL PATIENT IN AN AUTOMATED COLLECTION AND ANALYSIS PATIENT CARE SYSTEM


(57)

A system and a method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system are described. A set of collected measures retrieved on a substantially regular basis is periodically received from a medical device having a sensor for monitoring at least one physiological measure of an individual patient. The collected measures set includes individual measures which each relate to patient information recorded by the medical device. The collected measures set are stored into a patient care record for the individual patient within a database. The database is organized to store one or more patient care records which each include a plurality of the collected measures sets. Voice feedback is spoken by the individual patient into a remote client substantially contemporaneous to the collection of an identifiable device measures set. The voice feedback is processed into a set of quality of life measures which each relate to patient self assessment indicators. The identified collected device measures set and the quality of life measures set are received and stored into the patient care record for the individual patient within the database. The identified collected device measures set, the quality of life measures set, and one or more of the collected device measures sets in the patient care record for the individual patient are analyzed relative to one or more other collected device measures sets stored in the database server to determine a patient status indicator.





- (72) BARDY, GUST H., US
(71) BARDY, GUST H., US
(51) Int. Cl.⁷ G06F 17/60, G06F 17/30, H04L 12/16
(30) 1999/07/26 (09/361,777) US
(30) 1999/12/31 (09/476,600) US
(54) **SYSTEME ET METHODE DE NORMALISATION DE LA
RETROACTION VOCALE D'UN PATIENT DANS UN
SYSTEME AUTOMATISE DE COLLECTE ET D'ANALYSE DE
DONNEES SUR LES SOINS AUX PATIENTS**
(54) **SYSTEM AND METHOD FOR PROVIDING NORMALIZED
VOICE FEEDBACK FROM AN INDIVIDUAL PATIENT IN AN
AUTOMATED COLLECTION AND ANALYSIS PATIENT
CARE SYSTEM**



(57) A system and a method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system are described. A set of collected measures retrieved on a substantially regular basis is periodically received from a medical device having a sensor for monitoring at least one physiological measure of an individual patient. The collected measures set includes individual measures which each relate to patient information recorded by the medical device. The collected measures set are stored into a patient care record for the individual patient within a database. The database is organized to store one or more patient care records which each include a plurality of the collected measures sets. Voice feedback is spoken by the individual patient into a remote client substantially contemporaneous to the collection of an identifiable device measures set. The voice feedback is processed into a set of quality of life measures which each relate to patient self-assessment indicators. The identified collected device measures set and the quality of life measures set are received and stored into the patient care record for the individual patient within the database. The identified collected device measures set, the quality of life measures set, and one or more of the collected device measures sets in the patient care record for the individual patient are analyzed relative to one or more other collected device measures sets stored in the database server to determine a patient status indicator.

**SYSTEM AND METHOD FOR PROVIDING NORMALIZED VOICE
FEEDBACK FROM AN INDIVIDUAL PATIENT IN AN AUTOMATED
COLLECTION AND ANALYSIS PATIENT CARE SYSTEM**

Abstract

5 A system and a method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system are described. A set of collected measures retrieved on a substantially regular basis is periodically received from a medical device having a sensor for monitoring at least one physiological measure of an individual patient. The collected measures

10 set includes individual measures which each relate to patient information recorded by the medical device. The collected measures set are stored into a patient care record for the individual patient within a database. The database is organized to store one or more patient care records which each include a plurality of the collected measures sets. Voice feedback is spoken by the individual patient into a

15 remote client substantially contemporaneous to the collection of an identifiable device measures set. The voice feedback is processed into a set of quality of life measures which each relate to patient self-assessment indicators. The identified collected device measures set and the quality of life measures set are received and stored into the patient care record for the individual patient within the database.

20 The identified collected device measures set, the quality of life measures set, and one or more of the collected device measures sets in the patient care record for the individual patient are analyzed relative to one or more other collected device measures sets stored in the database server to determine a patient status indicator.

0128-01.ag3

5

SYSTEM AND METHOD FOR PROVIDING NORMALIZED VOICE
FEEDBACK FROM AN INDIVIDUAL PATIENT IN AN AUTOMATED
10 COLLECTION AND ANALYSIS PATIENT CARE SYSTEM

Cross-Reference to Related Application

This patent application is a continuation-in-part of commonly owned U.S. patent application, Serial No. 09/361,777, filed July 26, 1999, pending, which is a continuation-in-part of commonly owned U.S. patent application, Serial No.
15 09/324,894, pending, filed June 3, 1999, the disclosures of which are incorporated herein by reference.

Field of the Invention

The present invention relates in general to automated data collection and analysis, and, in particular, to a system and method for providing normalized
20 voice feedback from an individual patient in an automated collection and analysis patient care system.

Background of the Invention

A broad class of medical subspecialties, including cardiology, endocrinology, hematology, neurology, gastroenterology, urology,
25 ophthalmology, and otolaryngology, to name a few, rely on accurate and timely patient information for use in aiding health care providers in diagnosing and treating diseases and disorders. Often, proper medical diagnosis requires information on physiological events of short duration and sudden onset, yet these

types of events are often occur infrequently and with little or no warning. Fortunately, such patient information can be obtained via external, implantable, cutaneous, subcutaneous, and manual medical devices, and combinations thereof. For example, in the area of cardiology, implantable pulse generators (IPGs) are

5 medical devices commonly used to treat irregular heartbeats, known as arrhythmias. There are three basic types of IPGs. Cardiac pacemakers are used to manage bradycardia, an abnormally slow or irregular heartbeat. Bradycardia can cause symptoms such as fatigue, dizziness, and fainting. Implantable cardioverter

10 defibrillators (ICDs) are used to treat tachycardia, heart rhythms that are abnormally fast and life threatening. Tachycardia can result in sudden cardiac death (SCD). Finally, implantable cardiovascular monitors and therapeutic devices are used to monitor and treat structural problems of the heart, such as congestive heart failure, as well as rhythm problems.

Pacemakers and ICDs, as well as other types of implantable and external

15 medical devices, are equipped with an on-board, volatile memory in which telemetered signals can be stored for later retrieval and analysis. In addition, a growing class of cardiac medical devices, including implantable heart failure monitors, implantable event monitors, cardiovascular monitors, and therapy devices, are being used to provide similar stored device information. These

20 devices are able to store more than thirty minutes of per heartbeat data. Typically, the telemetered signals can provide patient device information recorded on a per heartbeat, binned average basis, or derived basis from, for example, atrial electrical activity, ventricular electrical activity, minute ventilation, patient activity score, cardiac output score, mixed venous oxygen score, cardiovascular

25 pressure measures, time of day, and any interventions and the relative success of such interventions. In addition, many such devices can have multiple sensors, or several devices can work together, for monitoring different sites within a patient's body.

Presently, stored device information is retrieved using a proprietary

30 interrogator or programmer, often during a clinic visit or following a device event.

The volume of data retrieved from a single device interrogation "snapshot" can be large and proper interpretation and analysis can require significant physician time and detailed subspecialty knowledge, particularly by cardiologists and cardiac electrophysiologists. The sequential logging and analysis of regularly scheduled
5 interrogations can create an opportunity for recognizing subtle and incremental changes in patient condition otherwise undetectable by inspection of a single "snapshot." However, present approaches to data interpretation and understanding and practical limitations on time and physician availability make such analysis impracticable.

10 Similarly, the determination and analysis of the quality of life issues which typically accompany the onset of a chronic yet stable diseases, such as coronary-artery disease, is a crucial adjunct to assessing patient wellness and progress. However, unlike in a traditional clinical setting, physicians participating in providing remote patient care are not able to interact with their patients in person.
15 Consequently, quality of life measures, such as how the patient subjectively looks and feels, whether the patient has shortness of breath, can work, can sleep, is depressed, is sexually active, can perform activities of daily life, and so on, cannot be implicitly gathered and evaluated.

A prior art system for collecting and analyzing pacemaker and ICD
20 telemetered signals in a clinical or office setting is the Model 9790 Programmer, manufactured by Medtronic, Inc., Minneapolis, MN. This programmer can be used to retrieve data, such as patient electrocardiogram and any measured physiological conditions, collected by the IPG for recordation, display and printing. The retrieved data is displayed in chronological order and analyzed by a
25 physician. Comparable prior art systems are available from other IPG manufacturers, such as the Model 2901 Programmer Recorder Monitor, manufactured by Guidant Corporation, Indianapolis, IN, which includes a removable floppy diskette mechanism for patient data storage. These prior art systems lack remote communications facilities and must be operated with the
30 patient present. These systems present a limited analysis of the collected data

based on a single device interrogation and lack the capability to recognize trends in the data spanning multiple episodes over time or relative to a disease specific peer group.

5 A prior art system for locating and communicating with a remote medical device implanted in an ambulatory patient is disclosed in U.S. Patent No. 5,752,976 ('976). The implanted device includes a telemetry transceiver for communicating data and operating instructions between the implanted device and an external patient communications device. The communications device includes a communication link to a remote medical support network, a global positioning
10 satellite receiver, and a patient activated link for permitting patient initiated communication with the medical support network. Patient voice communications through the patient link include both actual patient voice and manually actuated signaling which may convey an emergency situation. The patient voice is converted to an audio signal, digitized, encoded, and transmitted by data bus to a
15 system controller.

Related prior art systems for remotely communicating with and receiving telemetered signals from a medical device are disclosed in U.S. Patent Nos. 5,113,869 ('869) and 5,336,245 ('245). In the '869 patent, an implanted AECG monitor can be automatically interrogated at preset times of day to telemeter out
20 accumulated data to a telephonic communicator or a full disclosure recorder. The communicator can be automatically triggered to establish a telephonic communication link and transmit the accumulated data to an office or clinic through a modem. In the '245 patent, telemetered data is downloaded to a larger capacity, external data recorder and is forwarded to a clinic using an auto-dialer
25 and fax modem operating in a personal computer-based programmer/interrogator. However, the '976 telemetry transceiver, '869 communicator, and '245 programmer/interrogator are limited to facilitating communication and transferal of downloaded patient data and do not include an ability to automatically track, recognize, and analyze trends in the data itself. Moreover, the '976 telemetry
30 transceiver facilitates patient voice communications through transmission of a

digitized audio signal and does not perform voice recognition or other processing to the patient's voice.

In addition, the uses of multiple sensors situated within a patient's body at multiple sites are disclosed in U.S. Patent No. 5,040,536 ('536) and U.S. Patent 5 5,987,352 ('352). In the '536 patent, an intravascular pressure posture detector includes at least two pressure sensors implanted in different places in the cardiovascular system, such that differences in pressure with changes in posture are differentially measurable. However, the physiological measurements are used locally within the device, or in conjunction with any implantable device, to effect 10 a therapeutic treatment. In the '352 patent, an event monitor can include additional sensors for monitoring and recording physiological signals during arrhythmia and syncopal events. The recorded signals can be used for diagnosis, research or therapeutic study, although no systematic approach to analyzing these signals, particularly with respect to peer and general population groups, is 15 presented.

Thus, there is a need for a system and method for providing continuous retrieval, transferal, and automated analysis of retrieved medical device information, such as telemetered signals, retrieved in general from a broad class of implantable and external medical devices. Preferably, the automated analysis 20 would include recognizing a trend indicating disease onset, progression, regression, and status quo and determining whether medical intervention is necessary.

There is a further need for a system and method that would allow consideration of sets of collected measures, both actual and derived, from 25 multiple device interrogations. These collected measures sets could then be compared and analyzed against short and long term periods of observation.

There is a further need for a system and method that would enable the measures sets for an individual patient to be self-referenced and cross-referenced to similar or dissimilar patients and to the general patient population. Preferably, 30 the historical collected measures sets of an individual patient could be compared

and analyzed against those of other patients in general or of a disease specific peer group in particular.

There is a further need for a system and method for accepting and normalizing live voice feedback spoken by an individual patient while an identifiable set of telemetered signals is collected by a implantable medical device. Preferably, the normalized voice feedback a semi-quantitative self-assessment of an individual patient's physical and emotional well being at a time substantially contemporaneous to the collection of the telemetered signals.

Summary of the Invention

The present invention provides a system and method for automated collection and analysis of patient information retrieved from an implantable medical device for remote patient care. The patient device information relates to individual measures recorded by and retrieved from implantable medical devices, such as IPGs and monitors. The patient device information is received on a regular, e.g., daily, basis as sets of collected measures which are stored along with other patient records in a database. The information can be analyzed in an automated fashion and feedback provided to the patient at any time and in any location.

The present invention also provides a system and method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system. As before, patient device information is received on a regular, e.g., daily, basis as sets of collected measures which are stored along with other patient records in a database. Voice feedback spoken by an individual patient is processed into a set of quality of life measures by a remote client substantially contemporaneous to the recordation of an identifiable set of collected device measures by the implantable medical device. The processed voice feedback and identifiable collected device measures set are both received and stored into the patient record in the database for subsequent evaluation.

An embodiment of the present invention is a system and method for providing normalized voice feedback from an individual patient in an automated

collection and analysis patient care system. A set of device measures from a medical device adapted to be implanted is collected. The collected device measures set includes individual device measures which each relate to patient information recorded by the medical device adapted to be implanted for the individual patient. The collected device measures set from the medical device adapted to be implanted are periodically received over a communications link which is interfaced to a network server. The collected device measures set are stored into a patient care record for the individual patient within a database server. The database server is organized to store one or more patient care records which each include a plurality of the collected device measures sets. Voice feedback is spoken by the individual patient into a remote client substantially contemporaneous to the collection of an identifiable device measures set. The voice feedback is processed into a set of quality of life measures which each relate to patient self-assessment indicators. The identified collected device measures set and the quality of life measures set are received over the communications link interfaced to the network server respectively from the medical device adapted to be implanted and the remote client. The identified collected device measures set and the quality of life measures set are stored into the patient care record for the individual patient within the database server. The identified collected device measures set, the quality of life measures set, and one or more of the collected device measures sets in the patient care record for the individual patient are analyzed relative to one or more other collected device measures sets stored in the database server to determine a patient status indicator.

A further embodiment of the present invention is a system and method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system. A set of collected measures retrieved on a substantially regular basis is periodically received from a medical device having a sensor for monitoring at least one physiological measure of an individual patient. The collected measures set includes individual measures which each relate to patient information recorded by the medical device. The collected

measures set are stored into a patient care record for the individual patient within a database. The database is organized to store one or more patient care records which each include a plurality of the collected measures sets. Voice feedback is spoken by the individual patient into a remote client, which can include the
5 medical device itself, whether implantable, external or otherwise, substantially contemporaneous to the collection of an identifiable device measures set. The voice feedback is processed into a set of quality of life measures which each relate to patient self-assessment indicators. The identified collected device measures set and the quality of life measures set are received and stored into the patient care
10 record for the individual patient within the database. The identified collected device measures set, the quality of life measures set, and one or more of the collected device measures sets in the patient care record for the individual patient are analyzed relative to one or more other collected device measures sets stored in the database server to determine a patient status indicator.

15 The present invention facilitates the gathering, storage, and analysis of critical patient information obtained on a routine basis and analyzed in an automated manner. Thus, the burden on physicians and trained personnel to evaluate the volumes of information is significantly minimized while the benefits to patients are greatly enhanced.

20 The present invention also enables the simultaneous collection of both physiological measures from implantable medical devices and quality of life measures spoken in the patient's own words. Voice recognition technology enables the spoken patient feedback to be normalized to a standardized set of semi-quantitative quality of life measures, thereby facilitating holistic remote,
25 automated patient care.

Still other embodiments of the present invention will become readily apparent to those skilled in the art from the following detailed description, wherein is described embodiments of the invention by way of illustrating the best mode contemplated for carrying out the invention. As will be realized, the
30 invention is capable of other and different embodiments and its several details are

capable of modifications in various obvious respects, all without departing from the spirit and the scope of the present invention. Accordingly, the drawings and detailed description are to be regarded as illustrative in nature and not as restrictive.

5

Brief Description of the Drawings

FIGURE 1 is a block diagram showing a system for automated collection and analysis of patient information retrieved from an implantable medical device for remote patient care in accordance with the present invention;

FIGURE 2 is a block diagram showing the hardware components of the server system of the system of FIGURE 1;

FIGURE 3 is a block diagram showing the software modules of the server system of the system of FIGURE 1;

FIGURE 4 is a block diagram showing the analysis module of the server system of FIGURE 3;

FIGURE 5 is a database schema showing, by way of example, the organization of a cardiac patient care record stored in the database of the system of FIGURE 1;

FIGURE 6 is a record view showing, by way of example, a set of partial cardiac patient care records stored in the database of the system of FIGURE 1;

FIGURE 7 is a flow diagram showing a method for automated collection and analysis of patient information retrieved from an implantable medical device for remote patient care in accordance with the present invention;

FIGURE 8 is a flow diagram showing a routine for analyzing collected measures sets for use in the method of FIGURE 7;

FIGURE 9 is a flow diagram showing a routine for comparing sibling collected measures sets for use in the routine of FIGURE 8;

FIGURES 10A and 10B are flow diagrams showing a routine for comparing peer collected measures sets for use in the routine of FIGURE 8;

FIGURE 11 is a flow diagram showing a routine for providing feedback for use in the method of FIGURE 7;

FIGURE 12 is a block diagram showing a system for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system;

5 FIGURE 13 is a block diagram showing the software modules of the remote client of the system of FIGURE 12;

FIGURE 14 is a block diagram showing the software modules of the server system of the system of FIGURE 12;

10 FIGURE 15 is a database schema showing, by way of example, the organization of a quality of life record for cardiac patient care stored as part of a patient care record in the database of the system of FIGURE 12;

FIGURES 16A-16B are flow diagrams showing a method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system;

15 FIGURE 17 is a flow diagram showing a routine for processing voice feedback for use in the method of FIGURES 16A-16B;

FIGURE 18 is a flow diagram showing a routine for requesting a quality of life measure for use in the routine of FIGURE 17;

FIGURE 19 is a flow diagram showing a routine for recognizing and translating individual spoken words for use in the routine of FIGURE 17;

20 FIGURE 20 is a block diagram showing the software modules of the server system in a further embodiment of the system of FIGURE 12;

25 FIGURE 21 is a block diagram showing a system for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system in accordance with a further embodiment of the present invention;

FIGURE 22 is a block diagram showing the analysis module of the server system of FIGURE 21;

FIGURE 23 is a database schema showing, by way of example, the organization of a quality of life and symptom measures set record for care of

patients stored as part of a patient care record in the database of the system of
FIGURE 21;

FIGURE 24 is a record view showing, by way of example, a set of partial
cardiac patient care records stored in the database of the system of FIGURE 21;

5 FIGURE 25 is a Venn diagram showing, by way of example, peer group
overlap between the partial patient care records of FIGURE 24; and

FIGURES 26A-26B are flow diagrams showing a method for providing
normalized voice feedback from an individual patient in an automated collection
and analysis patient care system in accordance with a further embodiment of the
10 present invention.

Detailed Description

FIGURE 1 is a block diagram showing a system 10 for automated
collection and analysis of patient information retrieved from an implantable
medical device for remote patient care in accordance with the present invention.

15 A patient 11 is a recipient of an implantable medical device 12, such as, by way of
example, an IPG or a heart failure or event monitor, with a set of leads extending
into his or her heart. The implantable medical device 12 includes circuitry for
recording into a short-term, volatile memory telemetered signals, which are stored
as a set of collected measures for later retrieval.

20 For an exemplary cardiac implantable medical device, the telemetered
signals non-exclusively present patient information recorded on a per heartbeat,
binned average or derived basis and relating to: atrial electrical activity,
ventricular electrical activity, minute ventilation, patient activity score, cardiac
output score, mixed venous oxygenation score, cardiovascular pressure measures,
25 time of day, the number and types of interventions made, and the relative success
of any interventions, plus the status of the batteries and programmed settings.
Examples of pacemakers suitable for use in the present invention include the
Discovery line of pacemakers, manufactured by Guidant Corporation,
Indianapolis, IN. Examples of ICDs suitable for use in the present invention

include the Gem line of ICDs, manufactured by Medtronic Corporation, Minneapolis, MN.

In the described embodiment, the patient 11 has a cardiac implantable medical device. However, a wide range of related implantable medical devices are used in other areas of medicine and a growing number of these devices are also capable of measuring and recording patient information for later retrieval. These implantable medical devices include monitoring and therapeutic devices for use in metabolism, endocrinology, hematology, neurology, muscular disorders, gastroenterology, urology, ophthalmology, otolaryngology, orthopedics, and similar medical subspecialties. One skilled in the art would readily recognize the applicability of the present invention to these related implantable medical devices.

On a regular basis, the telemetered signals stored in the implantable medical device 12 are retrieved. By way of example, a programmer 14 can be used to retrieve the telemetered signals. However, any form of programmer, interrogator, recorder, monitor, or telemetered signals transceiver suitable for communicating with an implantable medical device 12 could be used, as is known in the art. In addition, a personal computer or digital data processor could be interfaced to the implantable medical device 12, either directly or via a telemetered signals transceiver configured to communicate with the implantable medical device 12.

Using the programmer 14, a magnetized reed switch (not shown) within the implantable medical device 12 closes in response to the placement of a wand 13 over the location of the implantable medical device 12. The programmer 14 communicates with the implantable medical device 12 via RF signals exchanged through the wand 13. Programming or interrogating instructions are sent to the implantable medical device 12 and the stored telemetered signals are downloaded into the programmer 14. Once downloaded, the telemetered signals are sent via an internetwork 15, such as the Internet, to a server system 16 which periodically receives and stores the telemetered signals in a database 17, as further described below with reference to FIGURE 2.

An example of a programmer 14 suitable for use in the present invention is the Model 2901 Programmer Recorder Monitor, manufactured by Guidant Corporation, Indianapolis, IN, which includes the capability to store retrieved telemetered signals on a proprietary removable floppy diskette. The telemetered signals could later be electronically transferred using a personal computer or similar processing device to the internetwork 15, as is known in the art.

Other alternate telemetered signals transfer means could also be employed. For instance, the stored telemetered signals could be retrieved from the implantable medical device 12 and electronically transferred to the internetwork 15 using the combination of a remote external programmer and analyzer and a remote telephonic communicator, such as described in U.S. Patent No. 5,113,869, the disclosure of which is incorporated herein by reference. Similarly, the stored telemetered signals could be retrieved and remotely downloaded to the server system 16 using a world-wide patient location and data telemetry system, such as described in U.S. Patent No. 5,752,976, the disclosure of which is incorporated herein by reference.

The received telemetered signals are analyzed by the server system 16, which generates a patient status indicator. The feedback is then provided back to the patient 11 through a variety of means. By way of example, the feedback can be sent as an electronic mail message generated automatically by the server system 16 for transmission over the internetwork 15. The electronic mail message is received by a remote client 18, such as a personal computer (PC), situated for local access by the patient 11. Alternatively, the feedback can be sent through a telephone interface device 19 as an automated voice mail message to a telephone 21 or as an automated facsimile message to a facsimile machine 22, both also situated for local access by the patient 11. In addition to a remote client 18, telephone 21, and facsimile machine 22, feedback could be sent to other related devices, including a network computer, wireless computer, personal data assistant, television, or digital data processor. Preferably, the feedback is

provided in a tiered fashion, as further described below with reference to FIGURE 3.

FIGURE 2 is a block diagram showing the hardware components of the server system 16 of the system 10 of FIGURE 1. The server system 16 consists of three individual servers: network server 31, database server 34, and application server 35. These servers are interconnected via an intranetwork 33. In the described embodiment, the functionality of the server system 16 is distributed among these three servers for efficiency and processing speed, although the functionality could also be performed by a single server or cluster of servers. The network server 31 is the primary interface of the server system 16 onto the internetwork 15. The network server 31 periodically receives the collected telemetered signals sent by remote implantable medical devices over the internetwork 15. The network server 31 is interfaced to the internetwork 15 through a router 32. To ensure reliable data exchange, the network server 31 implements a TCP/IP protocol stack, although other forms of network protocol stacks are suitable.

The database server 34 organizes the patient care records in the database 17 and provides storage of and access to information held in those records. A high volume of data in the form of collected measures sets from individual patients is received. The database server 34 frees the network server 31 from having to categorize and store the individual collected measures sets in the appropriate patient care record.

The application server 35 operates management applications and performs data analysis of the patient care records, as further described below with reference to FIGURE 3. The application server 35 communicates feedback to the individual patients either through electronic mail sent back over the internetwork 15 via the network server 31 or as automated voice mail or facsimile messages through the telephone interface device 19.

The server system 16 also includes a plurality of individual workstations 36 (WS) interconnected to the intranetwork 33, some of which can include

peripheral devices, such as a printer 37. The workstations 36 are for use by the data management and programming staff, nursing staff, office staff, and other consultants and authorized personnel.

The database 17 consists of a high-capacity storage medium configured to store individual patient care records and related health care information. Preferably, the database 17 is configured as a set of high-speed, high capacity hard drives, such as organized into a Redundant Array of Inexpensive Disks (RAID) volume. However, any form of volatile storage, non-volatile storage, removable storage, fixed storage, random access storage, sequential access storage, permanent storage, erasable storage, and the like would be equally suitable. The organization of the database 17 is further described below with reference to FIGURE 3.

The individual servers and workstations are general purpose, programmed digital computing devices consisting of a central processing unit (CPU), random access memory (RAM), non-volatile secondary storage, such as a hard drive or CD ROM drive, network interfaces, and peripheral devices, including user interfacing means, such as a keyboard and display. Program code, including software programs, and data are loaded into the RAM for execution and processing by the CPU and results are generated for display, output, transmittal, or storage. In the described embodiment, the individual servers are Intel Pentium-based server systems, such as available from Dell Computers, Austin, Texas, or Compaq Computers, Houston, Texas. Each system is preferably equipped with 128MB RAM, 100GB hard drive capacity, data backup facilities, and related hardware for interconnection to the intranetwork 33 and internetwork 15. In addition, the workstations 36 are also Intel Pentium-based personal computer or workstation systems, also available from Dell Computers, Austin, Texas, or Compaq Computers, Houston, Texas. Each workstation is preferably equipped with 64MB RAM, 10GB hard drive capacity, and related hardware for interconnection to the intranetwork 33. Other types of server and workstation systems, including personal computers, minicomputers, mainframe computers,

supercomputers, parallel computers, workstations, digital data processors and the like would be equally suitable, as is known in the art.

The telemetered signals are communicated over an internetwork 15, such as the Internet. However, any type of electronic communications link could be used, including an intranetwork link, serial link, data telephone link, satellite link, radio-frequency link, infrared link, fiber optic link, coaxial cable link, television link, and the like, as is known in the art. Also, the network server 31 is interfaced to the internetwork 15 using a T-1 network router 32, such as manufactured by Cisco Systems, Inc., San Jose, California. However, any type of interfacing device suitable for interconnecting a server to a network could be used, including a data modem, cable modem, network interface, serial connection, data port, hub, frame relay, digital PBX, and the like, as is known in the art.

FIGURE 3 is a block diagram showing the software modules of the server system 16 of the system 10 of FIGURE 1. Each module is a computer program written as source code in a conventional programming language, such as the C or Java programming languages, and is presented for execution by the CPU as object or byte code, as is known in the arts. The various implementations of the source code and object and byte codes can be held on a computer-readable storage medium or embodied on a transmission medium in a carrier wave. There are three basic software modules, which functionally define the primary operations performed by the server system 16: database module 51, analysis module 53, and feedback module 55. In the described embodiment, these modules are executed in a distributed computing environment, although a single server or a cluster of servers could also perform the functionality of the modules. The module functions are further described below in more detail beginning with reference to FIGURE 7.

For each patient being provided remote patient care, the server system 16 periodically receives a collected measures set 50 which is forwarded to the database module 51 for processing. The database module 51 organizes the individual patient care records stored in the database 52 and provides the facilities

for efficiently storing and accessing the collected measures sets 50 and patient data maintained in those records. An exemplary database schema for use in storing collected measures sets 50 in a patient care record is described below, by way of example, with reference to FIGURE 5. The database server 34 (shown in 5 FIGURE 2) performs the functionality of the database module 51. Any type of database organization could be utilized, including a flat file system, hierarchical database, relational database, or distributed database, such as provided by database vendors, such as Oracle Corporation, Redwood Shores, California.

The analysis module 53 analyzes the collected measures sets 50 stored in 10 the patient care records in the database 52. The analysis module 53 makes an automated determination of patient wellness in the form of a patient status indicator 54. Collected measures sets 50 are periodically received from implantable medical devices and maintained by the database module 51 in the database 52. Through the use of this collected information, the analysis module 15 53 can continuously follow the medical well being of a patient and can recognize any trends in the collected information that might warrant medical intervention. The analysis module 53 compares individual measures and derived measures obtained from both the care records for the individual patient and the care records for a disease specific group of patients or the patient population in general. The 20 analytic operations performed by the analysis module 53 are further described below with reference to FIGURE 4. The application server 35 (shown in FIGURE 2) performs the functionality of the analysis module 53.

The feedback module 55 provides automated feedback to the individual patient based, in part, on the patient status indicator 54. As described above, the 25 feedback could be by electronic mail or by automated voice mail or facsimile. Preferably, the feedback is provided in a tiered manner. In the described embodiment, four levels of automated feedback are provided. At a first level, an interpretation of the patient status indicator 54 is provided. At a second level, a notification of potential medical concern based on the patient status indicator 54 is 30 provided. This feedback level could also be coupled with human contact by

specially trained technicians or medical personnel. At a third level, the notification of potential medical concern is forwarded to medical practitioners located in the patient's geographic area. Finally, at a fourth level, a set of reprogramming instructions based on the patient status indicator 54 could be transmitted directly to the implantable medical device to modify the programming instructions contained therein. As is customary in the medical arts, the basic tiered feedback scheme would be modified in the event of bona fide medical emergency. The application server 35 (shown in FIGURE 2) performs the functionality of the feedback module 55.

FIGURE 4 is a block diagram showing the analysis module 53 of the server system 16 of FIGURE 3. The analysis module 53 contains two functional submodules: comparison module 62 and derivation module 63. The purpose of the comparison module 62 is to compare two or more individual measures, either collected or derived. The purpose of the derivation module 63 is to determine a derived measure based on one or more collected measures which is then used by the comparison module 62. For instance, a new and improved indicator of impending heart failure could be derived based on the exemplary cardiac collected measures set described with reference to FIGURE 5. The analysis module 53 can operate either in a batch mode of operation wherein patient status indicators are generated for a set of individual patients or in a dynamic mode wherein a patient status indicator is generated on the fly for an individual patient.

The comparison module 62 receives as inputs from the database 17 two input sets functionally defined as peer collected measures sets 60 and sibling collected measures sets 61, although in practice, the collected measures sets are stored on a per sampling basis. Peer collected measures sets 60 contain individual collected measures sets that all relate to the same type of patient information, for instance, atrial electrical activity, but which have been periodically collected over time. Sibling collected measures sets 61 contain individual collected measures sets that relate to different types of patient information, but which may have been collected at the same time or different times. In practice, the collected measures

sets are not separately stored as "peer" and "sibling" measures. Rather, each individual patient care record stores multiple sets of sibling collected measures. The distinction between peer collected measures sets 60 and sibling collected measures sets 61 is further described below with reference to FIGURE 6.

5 The derivation module 63 determines derived measures sets 64 on an as-needed basis in response to requests from the comparison module 62. The derived measures 64 are determined by performing linear and non-linear mathematical operations on selected peer measures 60 and sibling measures 61, as is known in the art.

10 FIGURE 5 is a database schema showing, by way of example, the organization of a cardiac patient care record stored 70 in the database 17 of the system 10 of FIGURE 1. Only the information pertaining to collected measures sets are shown. Each patient care record would also contain normal identifying and treatment profile information, as well as medical history and other pertinent data (not shown). Each patient care record stores a multitude of collected
15 measures sets for an individual patient. Each individual set represents a recorded snapshot of telemetered signals data which was recorded, for instance, per heartbeat or binned average basis by the implantable medical device 12. For example, for a cardiac patient, the following information would be recorded as a
20 collected measures set: atrial electrical activity 71, ventricular electrical activity 72, time of day 73, activity level 74, cardiac output 75, oxygen level 76, cardiovascular pressure measures 77, pulmonary measures 78, interventions made by the implantable medical device 78, and the relative success of any
25 interventions made 80. In addition, the implantable medical device 12 would also communicate device specific information, including battery status 81 and program settings 82. Other types of collected measures are possible. In addition, a well-documented set of derived measures can be determined based on the collected measures, as is known in the art.

30 FIGURE 6 is a record view showing, by way of example, a set of partial cardiac patient care records stored in the database 17 of the system 10 of FIGURE

1. Three patient care records are shown for *Patient 1*, *Patient 2*, and *Patient 3*. For each patient, three sets of measures are shown, *X*, *Y*, and *Z*. The measures are organized into sets with *Set 0* representing sibling measures made at a reference time $t=0$. Similarly, *Set n-2*, *Set n-1* and *Set n* each represent sibling measures
 5 made at later reference times $t=n-2$, $t=n-1$ and $t=n$, respectively.

For a given patient, for instance, *Patient 1*, all measures representing the same type of patient information, such as measure *X*, are peer measures. These are measures, which are monitored over time in a disease-matched peer group. All measures representing different types of patient information, such as measures
 10 *X*, *Y*, and *Z*, are sibling measures. These are measures which are also measured over time, but which might have medically significant meaning when compared to each other within a single set. Each of the measures, *X*, *Y*, and *Z*, could be either collected or derived measures.

The analysis module 53 (shown in FIGURE 4) performs two basic forms
 15 of comparison. First, individual measures for a given patient can be compared to other individual measures for that same patient. These comparisons might be peer-to-peer measures projected over time, for instance, $X_n, X_{n-1}, X_{n-2}, \dots, X_0$, or sibling-to-sibling measures for a single snapshot, for instance, X_n, Y_n , and Z_n , or projected over time, for instance, $X_n, Y_n, Z_n, X_{n-1}, Y_{n-1}, Z_{n-1}, X_{n-2}, Y_{n-2}, Z_{n-2}, \dots, X_0,$
 20 Y_0, Z_0 . Second, individual measures for a given patient can be compared to other individual measures for a group of other patients sharing the same disease-specific characteristics or to the patient population in general. Again, these comparisons might be peer-to-peer measures projected over time, for instance, $X_n,$
 $X_n', X_n'', X_{n-1}, X_{n-1}', X_{n-1}'', X_{n-2}, X_{n-2}', X_{n-2}'', \dots, X_0, X_0', X_0''$, or comparing the
 25 individual patient's measures to an average from the group. Similarly, these comparisons might be sibling-to-sibling measures for single snapshots, for instance, $X_n, X_n', X_n'', Y_n, Y_n', Y_n'',$ and Z_n, Z_n', Z_n'' , or projected over time, for instance, $X_n, X_n', X_n'', Y_n, Y_n', Y_n'', Z_n, Z_n', Z_n'', X_{n-1}, X_{n-1}', X_{n-1}'', Y_{n-1}, Y_{n-1}', Y_{n-1}'', Z_{n-1},$
 $Z_{n-1}', Z_{n-1}'', X_{n-2}, X_{n-2}', X_{n-2}'', Y_{n-2}, Y_{n-2}', Y_{n-2}'', Z_{n-2}, Z_{n-2}', Z_{n-2}'', \dots, X_0, X_0', X_0'', Y_0,$
 30 $Y_0', Y_0'',$ and Z_0, Z_0', Z_0'' . Other forms of comparisons are feasible.

FIGURE 7 is a flow diagram showing a method 90 for automated collection and analysis of patient information retrieved from an implantable medical device 12 for remote patient care in accordance with the present invention. The method 90 is implemented as a conventional computer program
5 for execution by the server system 16 (shown in FIGURE 1). As a preparatory step, the patient care records are organized in the database 17 with a unique patient care record assigned to each individual patient (block 91). Next, the collected measures sets for an individual patient are retrieved from the implantable medical device 12 (block 92) using a programmer, interrogator,
10 telemetered signals transceiver, and the like. The retrieved collected measures sets are sent, on a substantially regular basis, over the internetwork 15 or similar communications link (block 93) and periodically received by the server system 16 (block 94). The collected measures sets are stored into the patient care record in the database 17 for that individual patient (block 95). One or more of the
15 collected measures sets for that patient are analyzed (block 96), as further described below with reference to FIGURE 8. Finally, feedback based on the analysis is sent to that patient over the internetwork 15 as an email message, via telephone line as an automated voice mail or facsimile message, or by similar feedback communications link (block 97), as further described below with
20 reference to FIGURE 11.

FIGURE 8 is a flow diagram showing the routine for analyzing collected measures sets 96 for use in the method of FIGURE 7. The purpose of this routine is to make a determination of general patient wellness based on comparisons and heuristic trends analyses of the measures, both collected and derived, in the
25 patient care records in the database 17. A first collected measures set is selected from a patient care record in the database 17 (block 100). If the measures comparison is to be made to other measures originating from the patient care record for the same individual patient (block 101), a second collected measures set is selected from that patient care record (block 102). Otherwise, a group
30 measures comparison is being made (block 101) and a second collected measures

set is selected from another patient care record in the database 17 (block 103). Note the second collected measures set could also contain averaged measures for a group of disease specific patients or for the patient population in general.

Next, if a sibling measures comparison is to be made (block 104), a routine for comparing sibling collected measures sets is performed (block 105), as further described below with reference to FIGURE 9. Similarly, if a peer measures comparison is to be made (block 106), a routine for comparing sibling collected measures sets is performed (block 107), as further described below with reference to FIGURES 10A and 10B.

Finally, a patient status indicator is generated (block 108). By way of example, cardiac output could ordinarily be approximately 5.0 liters per minute with a standard deviation of ± 1.0 . An actionable medical phenomenon could occur when the cardiac output of a patient is ± 3.0 - 4.0 standard deviations out of the norm. A comparison of the cardiac output measures 75 (shown in FIGURE 5) for an individual patient against previous cardiac output measures 75 would establish the presence of any type of downward health trend as to the particular patient. A comparison of the cardiac output measures 75 of the particular patient to the cardiac output measures 75 of a group of patients would establish whether the patient is trending out of the norm. From this type of analysis, the analysis module 53 generates a patient status indicator 54 and other metrics of patient wellness, as is known in the art.

FIGURE 9 is a flow diagram showing the routine for comparing sibling collected measures sets 105 for use in the routine of FIGURE 8. Sibling measures originate from the patient care records for an individual patient. The purpose of this routine is either to compare sibling derived measures to sibling derived measures (blocks 111-113) or sibling collected measures to sibling collected measures (blocks 115-117). Thus, if derived measures are being compared (block 110), measures are selected from each collected measures set (block 111). First and second derived measures are derived from the selected measures (block 112) using the derivation module 63 (shown in FIGURE 4). The first and second

derived measures are then compared (block 113) using the comparison module 62 (also shown in FIGURE 4). The steps of selecting, determining, and comparing (blocks 111-113) are repeated until no further comparisons are required (block 114), whereupon the routine returns.

5 If collected measures are being compared (block 110), measures are selected from each collected measures set (block 115). The first and second collected measures are then compared (block 116) using the comparison module 62 (also shown in FIGURE 4). The steps of selecting and comparing (blocks 115-116) are repeated until no further comparisons are required (block 117),
10 whereupon the routine returns.

FIGURES 10A and 10B are a flow diagram showing the routine for comparing peer collected measures sets 107 for use in the routine of FIGURE 8. Peer measures originate from patient care records for different patients, including groups of disease specific patients or the patient population in general. The
15 purpose of this routine is to compare peer derived measures to peer derived measures (blocks 122-125), peer derived measures to peer collected measures (blocks 126-129), peer collected measures to peer derived measures (block 131-134), or peer collected measures to peer collected measures (blocks 135-137). Thus, if the first measure being compared is a derived measure (block 120) and
20 the second measure being compared is also a derived measure (block 121), measures are selected from each collected measures set (block 122). First and second derived measures are derived from the selected measures (block 123) using the derivation module 63 (shown in FIGURE 4). The first and second derived measures are then compared (block 124) using the comparison module 62
25 (also shown in FIGURE 4). The steps of selecting, determining, and comparing (blocks 122-124) are repeated until no further comparisons are required (block 115), whereupon the routine returns.

If the first measure being compared is a derived measure (block 120) but the second measure being compared is a collected measure (block 121), a first
30 measure is selected from the first collected measures set (block 126). A first

derived measure is derived from the first selected measure (block 127) using the derivation module 63 (shown in FIGURE 4). The first derived and second collected measures are then compared (block 128) using the comparison module 62 (also shown in FIGURE 4). The steps of selecting, determining, and comparing (blocks 126-128) are repeated until no further comparisons are required (block 129), whereupon the routine returns.

If the first measure being compared is a collected measure (block 120) but the second measure being compared is a derived measure (block 130), a second measure is selected from the second collected measures set (block 131). A second derived measure is derived from the second selected measure (block 132) using the derivation module 63 (shown in FIGURE 4). The first collected and second derived measures are then compared (block 133) using the comparison module 62 (also shown in FIGURE 4). The steps of selecting, determining, and comparing (blocks 131-133) are repeated until no further comparisons are required (block 134), whereupon the routine returns.

If the first measure being compared is a collected measure (block 120) and the second measure being compared is also a collected measure (block 130), measures are selected from each collected measures set (block 135). The first and second collected measures are then compared (block 136) using the comparison module 62 (also shown in FIGURE 4). The steps of selecting and comparing (blocks 135-136) are repeated until no further comparisons are required (block 137), whereupon the routine returns.

FIGURE 11 is a flow diagram showing the routine for providing feedback 97 for use in the method of FIGURE 7. The purpose of this routine is to provide tiered feedback based on the patient status indicator. Four levels of feedback are provided with increasing levels of patient involvement and medical care intervention. At a first level (block 150), an interpretation of the patient status indicator 54, preferably phrased in lay terminology, and related health care information is sent to the individual patient (block 151) using the feedback module 55 (shown in FIGURE 3). At a second level (block 152), a notification of

potential medical concern, based on the analysis and heuristic trends analysis, is sent to the individual patient (block 153) using the feedback module 55. At a third level (block 154), the notification of potential medical concern is forwarded to the physician responsible for the individual patient or similar health care professionals (block 155) using the feedback module 55. Finally, at a fourth level (block 156), reprogramming instructions are sent to the implantable medical device 12 (block 157) using the feedback module 55.

FIGURE 12 is a block diagram showing a system 200 for providing normalized voice feedback from an individual patient 11 in an automated collection and analysis patient care system, such as the system 10 of FIGURE 1. The remote client 18 includes a microphone 201 and a speaker 202 which is interfaced internally within the remote client 18 to sound recordation and reproduction hardware. The patient 11 provides spoken feedback into the microphone 201 in response to voice prompts reproduced by the remote client 18 on the speaker 202, as further described below with reference to FIGURE 13. The raw spoken feedback is processed into a normalized set of quality of life measures which each relate to uniform self-assessment indicators, as further described below with reference to FIGURE 15. Alternatively, in a further embodiment of the system 200, the patient 11 can provide spoken feedback via a telephone network 203 using a standard telephone 203, including a conventional wired telephone or a wireless telephone, such as a cellular telephone, as further described below with reference to FIGURE 20. In the described embodiment, the microphone 201 and the speaker 202 are standard, off-the-shelf components commonly included with consumer personal computer systems, as is known in the art.

The system 200 continuously monitors and collects sets of device measures from the implantable medical device 12. To augment the on-going monitoring process with a patient's self-assessment of physical and emotional well-being, a quality of life measures set can be recorded by the remote client 18. Importantly, each quality of life measures set is recorded substantially

contemporaneous to the collection of an identified collected device measures set .
The date and time of day at which the quality of life measures set was recorded
can be used to correlate the quality of life measures set to the collected device
measures set recorded closest in time to the quality of life measures set. The
5 pairing of the quality of life measures set and an identified collected device
measures set provides medical practitioners with a more complete picture of the
patient's medical status by combining physiological "hard" machine-recorded
data with semi-quantitative "soft" patient-provided data.

FIGURE 13 is a block diagram showing the software modules of the
10 remote client 18 of the system 200 of FIGURE 12. As with the software modules
of the system 10 of FIGURE 1, each module here is also a computer program
written as source code in a conventional programming language, such as the C or
Java programming languages, and is presented for execution by the CPU as object
or byte code, as is known in the arts. There are two basic software modules,
15 which functionally define the primary operations performed by the remote client
18 in providing normalized voice feedback: audio prompter 210 and speech
engine 214. The remote client 18 includes a secondary storage 219, such as a
hard drive, a CD ROM player, and the like, within which is stored data used by
the software modules. Conceptually, the voice reproduction and recognition
20 functions performed by the audio prompter 210 and speech engine 214 can be
described separately, but those same functions could also be performed by a
single voice processing module, as is known in the art.

The audio prompter 210 generates voice prompts 226 which are played
back to the patient 11 on the speaker 202. Each voice prompt is in the form of a
25 question or phrase seeking to develop a self-assessment of the patient's physical
and emotional well being. For example, the patient 11 might be prompted with,
"Are you short of breath?" The voice prompts 226 are either from a written script
220 reproduced by speech synthesizer 211 or pre-recorded speech 221 played
back by playback module 212. The written script 220 is stored within the
30 secondary storage 219 and consists of written quality of life measure requests.

Similarly, the pre-recorded speech 221 is also stored within the secondary storage 219 and consists of sound "bites" of recorded quality of life measure requests in either analog or digital format.

The speech engine 214 receives voice responses 227 spoken by the patient 5 11 into the microphone 201. The voice responses 227 can be unstructured, natural language phrases and sentences. A voice grammar 222 provides a lexical structuring for use in determining the meaning of each spoken voice response 227. The voice grammar 222 allows the speech engine 214 to "normalize" the voice responses 227 into recognized quality of life measures 228. Individual 10 spoken words in each voice response 227 are recognized by a speech recognition module 215 and translated into written words. In turn, the written words are parsed into tokens by a parser 216. A lexical analyzer 217 analyzes the tokens as complete phrases in accordance with a voice grammar 222 stored within the secondary storage 219. Finally, if necessary, the individual words are normalized 15 to uniform terms by a lookup module 218 which retrieves synonyms maintained as a vocabulary 223 stored within the secondary storage 218. For example, in response to the query, "Are you short of breath?" a patient might reply, "I can hardly breath," "I am panting," or "I am breathless." The speech recognition module 215 would interpret these phrases to imply dyspnea with a corresponding 20 quality of life measure indicating an awareness by the patient of abnormal breathing. In the described embodiment, the voice reproduction and recognition functions can be performed by the various natural voice software programs licensed by Dragon Systems, Inc., Newton, MA. Alternatively, the written script 220, voice grammar 222, and vocabulary 223 could be expressed as a script 25 written in a voice page markup language for interpretation by a voice browser operating on the remote client 18. Two exemplary voice page description languages include the VoXML markup language, licensed by Motorola, Inc., Chicago, IL, and described at <http://www.voxml.com>, and the Voice eXtensible Markup Language (VXML), currently being jointly developed by AT&T, 30 Motorola, Lucent Technologies, and IBM, and described at

<http://www.vxmlforum.com>. The module functions are further described below in more detail beginning with reference to FIGURES 16A-16B.

FIGURE 14 is a block diagram showing the software modules of the server system 16 of the system 200 of FIGURE 12. The database module 51, previously described above with reference to FIGURE 3, also receives the collected quality of life measures set 228 from the remote client 18, which the database module 51 stores into the appropriate patient care record in the database 52. The date and time of day 236 (shown in FIGURE 15) of the quality of life measures set 228 is matched to the date and time of day 73 (shown in FIGURE 5) of the collected measures set 50 recorded closest in time to the quality of life measures set 228. The matching collected measures set 50 is identified in the patient care record and can be analyzed with the quality of life measures set 228 by the analysis module 53, such as described above with reference to FIGURE 8.

FIGURE 15 is a database schema showing, by way of example, the organization of a quality of life record 230 for cardiac patient care stored as part of a patient care record in the database 17 of the system 200 of FIGURE 12. A quality of life score is a semi-quantitative self-assessment of an individual patient's physical and emotional well being. Non-commercial, non-proprietary standardized automated quality of life scoring systems are readily available, such as provided by the Duke Activities Status Indicator. For example, for a cardiac patient, the quality of life record 230 stores the following information: health wellness 231, shortness of breath 232, energy level 233, chest discomfort 235, time of day 234, and other quality of life measures as would be known to one skilled in the art. Other types of quality of life measures are possible.

A quality of life indicator is a vehicle through which a patient can remotely communicate to the patient care system how he or she is subjectively feeling. The quality of life indicators can include symptoms of disease. When tied to machine-recorded physiological measures, a quality of life indicator can provide valuable additional information to medical practitioners and the automated collection and analysis patient care system 200 not otherwise

discernible without having the patient physically present. For instance, a scoring system using a scale of 1.0 to 10.0 could be used with 10.0 indicating normal wellness and 1.0 indicating severe health problems. Upon the completion of an initial observation period, a patient might indicate a health wellness score 231 of 5.0 and a cardiac output score of 5.0. After one month of remote patient care, the patient might then indicate a health wellness score 231 of 4.0 and a cardiac output score of 4.0 and a week later indicate a health wellness score 231 of 3.5 and a cardiac output score of 3.5. Based on a comparison of the health wellness scores 231 and the cardiac output scores, the system 200 would identify a trend indicating the necessity of potential medical intervention while a comparison of the cardiac output scores alone might not lead to the same prognosis.

FIGURES 16A-16B are flow diagrams showing a method 239 for providing normalized voice feedback from an individual patient 11 in an automated collection and analysis patient care system 200. As with the method 90 of FIGURE 7, this method is also implemented as a conventional computer program and performs the same set of steps as described with reference to FIGURE 7 with the following additional functionality. First, voice feedback spoken by the patient 11 into the remote client 18 is processed into a quality of life measures set 228 (block 240), as further described below with reference to FIGURE 17. The voice feedback is spoken substantially contemporaneous to the collection of an identified device measures set 50. The appropriate collected device measures set 50 can be matched to and identified with (not shown) the quality of life measures set 228 either by matching their respective dates and times of day or by similar means, either by the remote client 18 or the server system 16. The quality of life measures set 228 and the identified collected measures set 50 are sent over the internetwork 15 to the server system 16 (block 241). Note the quality of life measures set 228 and the identified collected measures set 50 both need not be sent over the internetwork 15 at the same time, so long as the two sets are ultimately paired based on, for example, date and time of day. The quality of life measures set 228 and the identified collected measures

set 50 are received by the server system 16 (block 242) and stored in the appropriate patient care record in the database 52 (block 243). Finally, the quality of life measures set 228, identified collected measures set 50, and one or more collected measures sets 50 are analyzed (block 244) and feedback, including a patient status indicator 54 (shown in FIGURE 14), is provided to the patient (block 245).

FIGURE 17 is a flow diagram showing the routine for processing voice feedback 240 for use in the method of FIGURES 16A-16B. The purpose of this routine is to facilitate a voice interactive session with the patient 11 during which is developed a normalized set of quality of life measures. Thus, the remote client 18 requests a quality of life measure via a voice prompt (block 250), played on the speaker 202 (shown in FIGURE 13), as further described below with reference to FIGURE 18. The remote client 18 receives the spoken feedback from the patient 11 (block 251) via the microphone 201 (shown in FIGURE 13). The remote client 18 recognizes individual words in the spoken feedback and translates those words into written words (block 252), as further described below with reference to FIGURE 19. The routine returns at the end of the voice interactive session.

FIGURE 18 is a flow diagram showing the routine for requesting a quality of life measure 251 for use in the routine 240 of FIGURE 17. The purpose of this routine is to present a voice prompt 226 to the user via the speaker 202. Either pre-recorded speech 221 or speech synthesized from a written script 220 can be used. Thus, if synthesized speech is employed by the remote client 18 (block 260), a written script, such as a voice markup language script, specifying questions and phrases which with to request quality of life measures is stored (block 261) on the secondary storage 219 of the remote client 18. Each written quality of life measure request is retrieved by the remote client 18 (block 262) and synthesized into speech for playback to the patient 11 (block 263). Alternatively, if pre-recorded speech is employed by the remote client 18 (block 260), pre-recorded voice "bites" are stored (block 264) on the secondary storage 219 of the remote client 18. Each pre-recorded quality of life measure request is retrieved by

the remote client 18 (block 265) and played back to the patient 11 (block 266).
The routine then returns.

FIGURE 19 is a flow diagram showing the routine for recognizing and translating individual spoken words 252 for use in the routine 240 of FIGURE 17.
5 The purpose of this routine is to receive and interpret a free-form voice response 227 from the user via the microphone 201. First, a voice grammar consisting of a lexical structuring of words, phrases, and sentences is stored (block 270) on the secondary storage 219 of the remote client 18. Similarly, a vocabulary of individual words and their commonly accepted synonyms is stored (block 271) on
10 the secondary storage 219 of the remote client 18. After individual words in the voice feedback are recognized (block 272), the individual words are parsed into tokens (block 273). The voice feedback is then lexically analyzed using the tokens and in accordance with the voice grammar 222 (block 274) to determine the meaning of the voice feedback. If necessary, the vocabulary 223 is referenced
15 to lookup synonyms of the individual words (block 275). The routine then returns.

FIGURE 20 is a block diagram showing the software modules of the server system in a further embodiment of the system 200 of FIGURE 12. The functionality of the remote client 18 in providing normalized voice feedback is
20 incorporated directly into the server system 16. The system 200 of FIGURE 12 requires the patient 11 to provide spoken feedback via a locally situated remote client 18. However, the system 280 enables a patient 11 to alternatively provide spoken feedback via a telephone network 203 using a standard telephone 203, including a conventional wired telephone or a wireless telephone, such as a
25 cellular telephone. The server system 16 is augmented to include the audio prompter 210, the speech engine 214, and the data stored in the secondary storage 219. A telephonic interface 280 interfaces the server system 16 to the telephone network 203 and receives voice responses 227 and sends voice prompts 226 to and from the server system 16. Telephonic interfacing devices are commonly
30 known in the art.

FIGURE 21 is a block diagram showing a system for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system 300 in accordance with a further embodiment of the present invention. The system 300 provides remote patient care in a manner similar to the system 200 of FIGURE 12, but with additional functionality for diagnosing and monitoring multiple sites within a patient's body using a variety of patient sensors for diagnosing one or more disorder. The patient 301 can be the recipient of an implantable medical device 302, as described above, or have an external medical device 303 attached, such as a Holter monitor-like device for monitoring electrocardiograms. In addition, one or more sites in or around the patient's body can be monitored using multiple sensors 304a, 304b, such as described in U.S. Patents 4,987,897; 5,040,536; 5,113,859; and 5,987,352, the disclosures of which are incorporated herein by reference. One automated system and method for collecting and analyzing retrieved patient information suitable for use with the present invention is described in the related, commonly-owned U.S. Patent application, Serial No. _____, entitled "System And Method For Automated Collection And Analysis Of Regularly Retrieved Patient Information For Remote Patient Care," pending, filed December 31, 1999, the disclosure of which is incorporated herein by reference. Other types of devices with physiological measure sensors, both heterogeneous and homogenous, could be used, either within the same device or working in conjunction with each other, as is known in the art.

As part of the system 300, the database 17 stores patient care records 305 for each individual patient to whom remote patient care is being provided. Each patient care record 305 contains normal patient identification and treatment profile information, as well as medical history, medications taken, height and weight, and other pertinent data (not shown). The patient care records 305 consist primarily of monitoring sets 306 storing device and derived measures (D&DM) sets 307 and quality of life and symptom measures (QOLM) sets 308 recorded and determined thereafter on a regular, continuous basis. The organization of the

device and derived measures sets 305 for an exemplary cardiac patient care record is described above with reference to FIGURE 5. The organization of the quality of life and symptom measures sets 308 is further described below with reference to FIGURE 23.

5 Optionally, the patient care records 305 can further include a reference baseline 309 storing a special set of device and derived reference measures sets 310 and quality of life and symptom measures sets 311 recorded and determined during an initial observation period, such as described in the related, commonly-owned U.S. Patent application, Serial No. _____, entitled "System And Method For
10 Determining A Reference Baseline Of Individual Patient Status For Use In An Automated Collection And Analysis Patient Care System," pending, filed December 31, 1999, the disclosure of which is incorporated herein by reference. Other forms of database organization are feasible.

 Finally, simultaneous notifications can also be delivered to the patient's
15 physician, hospital, or emergency medical services provider 312 using feedback means similar to that used to notify the patient. As described above, the feedback could be by electronic mail or by automated voice mail or facsimile. Furthermore, the spoken voice feedback from the patient and the feedback provided by the system 200 can be communicated by means of or in combination
20 with the medical device itself, whether implantable, external or otherwise.

 FIGURE 22 is a block diagram showing the analysis module 53 of the server system 16 of FIGURE 21. The peer collected measures sets 60 and sibling collected measures sets 61 can be organized into site specific groupings based on the sensor from which they originate, that is, implantable medical device 302,
25 external medical device 303, or multiple sensors 304a, 304b. The functionality of the analysis module 53 is augmented to iterate through a plurality of site specific measures sets 315 and one or more disorders.

 As described above, as an adjunct to remote patient care through the monitoring of measured physiological data via implantable medical device 302,
30 external medical device 303 and multiple sensors 304a, 304b, quality of life and

symptom measures sets 308 can also be stored in the database 17 as part of the monitoring sets 306. A quality of life measure is a semi-quantitative self-assessment of an individual patient's physical and emotional well-being and a record of symptoms, such as provided by the Duke Activities Status Indicator.

5 These scoring systems can be provided for use by the patient 11 on the personal computer 18 (shown in FIGURE 1) to record his or her quality of life scores for both initial and periodic download to the server system 16.

FIGURE 23 is a database schema which augments the database schema described above with reference to FIGURE 15 and showing, by way of example,
10 the organization of a quality of life and symptom measures set record 320 for care of patients stored as part of a patient care record 305 in the database 17 of the system 300 of FIGURE 21. The following exemplary information is recorded for a patient: overall health wellness 321, psychological state 322, chest discomfort 323, location of chest discomfort 324, palpitations 325, shortness of breath 326,
15 exercise tolerance 327, cough 328, sputum production 329, sputum color 330, energy level 331, syncope 332, near syncope 333, nausea 334, diaphoresis 335, time of day 91, and other quality of life and symptom measures as would be known to one skilled in the art.

Other types of quality of life and symptom measures are possible, such as
20 those indicated by responses to the Minnesota Living with Heart Failure Questionnaire described in E. Braunwald, ed., "Heart Disease—A Textbook of Cardiovascular Medicine," pp. 452-454, W.B. Saunders Co. (1997), the disclosure of which is incorporated herein by reference. Similarly, functional classifications based on the relationship between symptoms and the amount of effort required to
25 provoke them can serve as quality of life and symptom measures, such as the New York Heart Association (NYHA) classifications I, II, III and IV, also described in *Ibid*.

The patient may also add non-device quantitative measures, such as the six-minute walk distance, as complementary data to the device and derived

measures sets 307 and the symptoms during the six-minute walk to quality of life and symptom measures sets 308.

FIGURE 24 is a record view showing, by way of example, a set of partial cardiac patient care records stored in the database 17 of the system 300 of

5 FIGURE 21. Three patient care records are again shown for *Patient 1*, *Patient 2*, and *Patient 3* with each of these records containing site specific measures sets 315, grouped as follows. First, the patient care record for *Patient 1* includes three site specific measures sets *A*, *B* and *C*, corresponding to three sites on *Patient 1*'s body. Similarly, the patient care record for *Patient 2* includes two site specific

10 measures sets *A* and *B*, corresponding to two sites, both of which are in the same relative positions on *Patient 2*'s body as the sites for *Patient 1*. Finally, the patient care record for *Patient 3* includes two site specific measures sets *A* and *D*, also corresponding to two medical device sensors, only one of which, Site *A*, is in the same relative position as Site *A* for *Patient 1* and *Patient 2*.

15 The analysis module 53 (shown in FIGURE 22) performs two further forms of comparison in addition to comparing the individual measures for a given patient to other individual measures for that same patient or to other individual measures for a group of other patients sharing the same disease-specific characteristics or to the patient population in general. First, the individual

20 measures corresponding to each body site for an individual patient can be compared to other individual measures for that same patient, a peer group or a general patient population. Again, these comparisons might be peer-to-peer measures projected over time, for instance, comparing measures for each site, *A*,

B and *C*, for *Patient 1*, $X_{n_1}, X_{n_1^*}, X_{n_1^*}, X_{n-1_1}, X_{n-1_1^*}, X_{n-1_1^*}, X_{n-2_1}, X_{n-2_1^*},$

25 $X_{n-2_1^*}, \dots, X_{n_2}, X_{n_2^*}, X_{n_2^*}; X_{n_2}, X_{n_2^*}, X_{n_2^*}, X_{n-1_2}, X_{n-1_2^*}, X_{n-1_2^*}, X_{n-2_2}, X_{n-2_2^*},$
 $X_{n-2_2^*}, \dots, X_{n_3}, X_{n_3^*}, X_{n_3^*}; X_{n_3}, X_{n_3^*}, X_{n_3^*}, X_{n-1_3}, X_{n-1_3^*}, X_{n-1_3^*}, X_{n-2_3}, X_{n-2_3^*},$
 $X_{n-2_3^*}, \dots, X_{n_4}, X_{n_4^*}, X_{n_4^*};$ comparing comparable measures for Site *A* for the

three patients, $X_{n_1}, X_{n_1^*}, X_{n_1^*}, X_{n-1_1}, X_{n-1_1^*}, X_{n-1_1^*}, X_{n-2_1}, X_{n-2_1^*}, X_{n-2_1^*}, \dots,$

$X_{n_i}, X_{q_i}, X_{o_i}$; or comparing the individual patient's measures to an average from the group. Similarly, these comparisons might be sibling-to-sibling measures for single snapshots, for instance, comparing comparable measures for Site A for the three patients, $X_{n_i}, X_{n_j}, X_{n_k}, Y_{n_i}, Y_{n_j}, Y_{n_k}$, and $Z_{n_i}, Z_{n_j}, Z_{n_k}$,
 5 or comparing those same comparable measures for Site A projected over time, for instance, $X_{n_i}, X_{n_j}, X_{n_k}, Y_{n_i}, Y_{n_j}, Y_{n_k}, Z_{n_i}, Z_{n_j}, Z_{n_k}, X_{n-1_i}, X_{n-1_j}, X_{n-1_k},$
 $Y_{n-1_i}, Y_{n-1_j}, Y_{n-1_k}, Z_{n-1_i}, Z_{n-1_j}, Z_{n-1_k}, X_{n-2_i}, X_{n-2_j}, X_{n-2_k}, Y_{n-2_i}, Y_{n-2_j},$
 $Y_{n-2_k}, Z_{n-2_i}, Z_{n-2_j}, Z_{n-2_k}, \dots, X_{o_i}, X_{o_j}, X_{o_k}, Y_{o_i}, Y_{o_j}, Y_{o_k}$, and $Z_{o_i}, Z_{o_j}, Z_{o_k}$.
 Other forms of site-specific comparisons, including comparisons between
 10 individual measures from non-comparable sites between patients, are feasible.

Second, the individual measures can be compared on a disorder specific basis. The individual measures stored in each cardiac patient record can be logically grouped into measures relating to specific disorders and diseases, for instance, congestive heart failure, myocardial infarction, respiratory distress, and
 15 atrial fibrillation. The foregoing comparison operations performed by the analysis module 53 are further described below with reference to FIGURES 26A-26B.

FIGURE 25 is a Venn diagram showing, by way of example, peer group overlap between the partial patient care records 305 of FIGURE 24. Each patient care record 305 includes characteristics data 350, 351, 352, including personal
 20 traits, demographics, medical history, and related personal data, for patients 1, 2 and 3, respectively. For example, the characteristics data 350 for patient 1 might include personal traits which include gender and age, such as male and an age between 40-45; a demographic of resident of New York City; and a medical history consisting of anterior myocardial infarction, congestive heart failure and
 25 diabetes. Similarly, the characteristics data 351 for patient 2 might include identical personal traits, thereby resulting in partial overlap 353 of characteristics data 350 and 351. Similar characteristics overlap 354, 355, 356 can exist between each respective patient. The overall patient population 357 would include the

universe of all characteristics data. As the monitoring population grows, the number of patients with personal traits matching those of the monitored patient will grow, increasing the value of peer group referencing. Large peer groups, well matched across all monitored measures, will result in a well known natural
5 history of disease and will allow for more accurate prediction of the clinical course of the patient being monitored. If the population of patients is relatively small, only some traits 356 will be uniformly present in any particular peer group. Eventually, peer groups, for instance, composed of 100 or more patients each, would evolve under conditions in which there would be complete overlap of
10 substantially all salient data, thereby forming a powerful core reference group for any new patient being monitored.

FIGURES 26A-26B are flow diagrams showing a method for providing normalized voice feedback from an individual patient in an automated collection and analysis patient care system 360 in accordance with a further embodiment of
15 the present invention. As with the method 239 of FIGURES 16A and 16B, this method is also implemented as a conventional computer program and performs the same set of steps as described with reference to FIGURES 16A and 16B with the following additional functionality. As before, the patient care records are organized in the database 17 with a unique patient care record assigned to each
20 individual patient (block 361). Next, the individual measures for each site are iteratively obtained in a first processing loop (blocks 362-367) and each disorder is iteratively analyzed in a second processing loop (blocks 368-370). Other forms of flow control are feasible, including recursive processing.

During each iteration of the first processing loop (blocks 362-367), the
25 collected measures sets for an individual patient are retrieved from the medical device or sensor located at the current site (block 363) using a programmer, interrogator, telemetered signals transceiver, and the like. The retrieved collected measures sets are sent, on a substantially regular basis, over the internetwork 15 or similar communications link (block 364) and periodically received by the
30 server system 16 (block 365). The collected measures sets are stored into the

patient care record 305 in the database 17 for that individual patient (block 366). Any voice feedback spoken by the patient 11 into the remote client 18 is processed into a quality of life measures set 228 (block 240), as described above with reference to FIGURE 17. The voice feedback is spoken substantially
5 contemporaneous to the collection of an identified device measures set 50. The appropriate collected device measures set 50 can be matched to and identified with (not shown) the quality of life measures set 228 either by matching their respective dates and times of day or by similar means, either by the remote client 18 or the server system 16. The quality of life measures set 228 and the identified
10 collected measures set 50 are sent over the internetwork 15 to the server system 16 (block 241). The quality of life measures set 228 and the identified collected measures set 50 are received by the server system 16 (block 242) and stored in the appropriate patient care record in the database 52 (block 243).

During each iteration of the second processing loop (blocks 368-370), the
15 quality of life measures set 228, identified collected measures set 50, and one or more of the collected measures sets for that patient are analyzed for the current disorder are analyzed (block 244). Finally, feedback based on the analysis is sent to that patient over the internetwork 15 as an email message, via telephone line as an automated voice mail or facsimile message, or by similar feedback
20 communications link (block 245). In addition, the measures sets can be further evaluated and matched to diagnose specific medical disorders, such as congestive heart failure, myocardial infarction, respiratory distress, and atrial fibrillation, as described in related, commonly-owned U.S. Patent applications, Serial No. 09/441,623, pending, filed November 16, 1999; Serial No. 09/441,612, pending,
25 filed November 16, 1999; Serial No. 09/442, 125, pending, filed November 16, 1999; and Serial No. 09/441,613, pending, filed November 16, 1999, the disclosures of which are incorporated herein by reference. In addition, multiple near-simultaneous disorders can be ordered and prioritized as part of the patient status indicator as described in the related, commonly-owned U.S. Patent

application, Serial No. 09/441,405, pending, filed November 16, 1999, the disclosure of which is incorporated herein by reference.

Therefore, through the use of the collected measures sets, the present invention makes possible immediate access to expert medical care at any time and in any place. For example, after establishing and registering for each patient an appropriate baseline set of measures, the database server could contain a virtually up-to-date patient history, which is available to medical providers for the remote diagnosis and prevention of serious illness regardless of the relative location of the patient or time of day.

Moreover, the gathering and storage of multiple sets of critical patient information obtained on a routine basis makes possible treatment methodologies based on an algorithmic analysis of the collected data sets. Each successive introduction of a new collected measures set into the database server would help to continually improve the accuracy and effectiveness of the algorithms used. In addition, the present invention potentially enables the detection, prevention, and cure of previously unknown forms of disorders based on a trends analysis and by a cross-referencing approach to create continuously improving peer-group reference databases.

Similarly, the present invention makes possible the provision of tiered patient feedback based on the automated analysis of the collected measures sets. This type of feedback system is suitable for use in, for example, a subscription based health care service. At a basic level, informational feedback can be provided by way of a simple interpretation of the collected data. The feedback could be built up to provide a graduated response to the patient, for example, to notify the patient that he or she is trending into a potential trouble zone. Human interaction could be introduced, both by remotely situated and local medical practitioners. Finally, the feedback could include direct interventive measures, such as remotely reprogramming a patient's IPG.

Finally, the present invention allows "live" patient voice feedback to be captured simultaneously with the collection of physiological measures by their

implantable medical device. The voice feedback is normalized to a standardized set of quality of life measures which can be analyzed in a remote, automated fashion. The voice feedback could also be coupled with visual feedback, such as through digital photography or video, to provide a more complete picture of the patient's physical well-being.

While the invention has been particularly shown and described as referenced to the embodiments thereof, those skilled in the art will understand that the foregoing and other changes in form and detail may be made therein without departing from the spirit and scope of the invention.

What is claimed is:

- 1 1. A system for providing normalized voice feedback from an
 2 individual patient in an automated collection and analysis patient care system,
 3 comprising:
 4 a medical device having a sensor for monitoring at least one physiological
 5 measure of an individual patient, the collected measures set comprising individual
 6 measures which each relate to patient information recorded by the medical device;
 7 a remote client processing voice feedback into a set of quality of life
 8 measures which each relate to patient self-assessment indicators, the voice
 9 feedback having been spoken by the individual patient into a remote client
 10 substantially contemporaneous to the collection of an identifiable device measures
 11 set;
 12 a network server over which is periodically received a set of collected
 13 measures retrieved on a substantially regular basis from the medical device, the
 14 identified collected device measures set is received from the medical device, and
 15 the quality of life measures set is received from the remote client;
 16 a database coupled to the network server and storing the collected
 17 measures set, the identified collected device measures set and the quality of life
 18 measures set into a patient care record for the individual patient within a database
 19 organized to store one or more patient care records which each comprise a
 20 plurality of the collected measures sets; and
 21 an application server coupled to the database and analyzing the identified
 22 collected device measures set, the quality of life measures set, and one or more of
 23 the collected device measures sets in the patient care record for the individual
 24 patient relative to one or more other collected device measures sets stored in the
 25 database to determine a patient status indicator.
- 1 2. A system according to Claim 1, further comprising:
 2 the network server repeatedly receiving one or more collected measures
 3 sets which are each recorded by a sensor which monitors at least one

4 physiological measure of the individual patient, each such sensor monitoring a
5 site within the individual patient unique from the site monitored by any other such
6 sensor;
7 the database storing each collected measures set organized by specific site
8 into the patient care record for the individual patient within the database; and
9 the application server analyzing one or more of the site specific collected
10 measures sets in the patient care record for each site within the individual patient
11 relative to one or more other site specific collected measures sets stored in the
12 database to determine a patient status indicator.

1 3. A system according to Claim 2, wherein the one or more site
2 specific collected measures sets and the one or more other site specific collected
3 measures sets both store measures collected from the same relative site.

1 4. A system according to Claim 2, wherein the one or more site
2 specific collected measures sets and the one or more other site specific collected
3 measures sets both store measures collected from a different site.

1 5. A system according to Claim 1, the remote client further
2 comprising:
3 an audio prompter requesting a quality of life measure via a voice prompt
4 played on the remote client to the individual patient.

1 6. A system according to Claim 5, further comprising:
2 a written script comprising a plurality of quality of life measure requests
3 stored within the remote client; and
4 the audio prompter further comprising a speech synthesizer module
5 retrieving each quality of life request from the stored written script with each such
6 retrieved quality of life measure request comprising one such voice prompt and
7 synthesizing speech for playback from the retrieved quality of life request.

1 7. A system according to Claim 5, further comprising:

2 pre-recorded speech comprising a plurality of quality of life measure
3 requests stored within the remote client; and
4 the audio prompter further comprising a playback module retrieving each
5 quality of life request from the stored pre-recorded speech with each such
6 retrieved quality of life measure request comprising one such voice prompt and
7 playing the pre-recorded speech from the retrieved quality of life request.

1 8. A system according to Claim 1, the remote client further
2 comprising:

3 a speech engine recognizing individual words in the spoken voice
4 feedback and translating the individual spoken words into written individual
5 words.

1 9. A system according to Claim 8, further comprising:

2 a voice grammar stored within the remote client, the voice grammar
3 comprising a plurality of speech phrases expressed in a natural language, each
4 speech phrase corresponding to a normalized quality of life measure;

5 the speech engine further comprising:

6 a parser parsing the written individual words into tokens; and
7 a lexical analyzer performing a lexical analysis of the parsed
8 tokens in accordance with the voice grammar to identify one such normalized
9 quality of life measure.

1 10. A system according to Claim 8, further comprising:

2 a vocabulary stored within the remote client, the vocabulary comprising
3 the written individual words; and

4 the speech engine further comprising a lookup module performing
5 a lookup of the written individual words from the vocabulary stored within the
6 remote client.

1 11. A system according to Claim 1, the remote client further
2 comprising:

3 wherein the remote client comprises at least one of a personal computer,
4 an audio interface, and a telephony instrument.

1 12. A method for providing normalized voice feedback from an
2 individual patient in an automated collection and analysis patient care system,
3 comprising:
4 periodically receiving a set of collected measures retrieved on a
5 substantially regular basis from a medical device having a sensor for monitoring
6 at least one physiological measure of an individual patient, the collected measures
7 set comprising individual measures which each relate to patient information
8 recorded by the medical device;
9 storing the collected measures set into a patient care record for the
10 individual patient within a database organized to store one or more patient care
11 records which each comprise a plurality of the collected measures sets;
12 processing voice feedback into a set of quality of life measures which each
13 relate to patient self-assessment indicators, the voice feedback having been
14 spoken by the individual patient into a remote client substantially
15 contemporaneous to the collection of an identifiable device measures set;
16 receiving the identified collected device measures set and the quality of
17 life measures set;
18 storing the identified collected device measures set and the quality of life
19 measures set into the patient care record for the individual patient within the
20 database; and
21 analyzing the identified collected device measures set, the quality of life
22 measures set, and one or more of the collected device measures sets in the patient
23 care record for the individual patient relative to one or more other collected
24 device measures sets stored in the database to determine a patient status indicator.

1 13. A method according to Claim 12, further comprising:
2 repeatedly receiving one or more collected measures sets which are each
3 recorded by a sensor which monitors at least one physiological measure of the

4 individual patient, each such sensor monitoring a site within the individual patient
5 unique from the site monitored by any other such sensor;
6 storing each collected measures set organized by specific site into the
7 patient care record for the individual patient within the database; and
8 analyzing one or more of the site specific collected measures sets in the
9 patient care record for each site within the individual patient relative to one or
10 more other site specific collected measures sets stored in the database to
11 determine a patient status indicator.

1 14. A method according to Claim 13, wherein the one or more site
2 specific collected measures sets and the one or more other site specific collected
3 measures sets both store measures collected from the same relative site.

1 15. A method according to Claim 13, wherein the one or more site
2 specific collected measures sets and the one or more other site specific collected
3 measures sets both store measures collected from a different site.

1 16. A method according to Claim 12, the operation of processing voice
2 feedback further comprising:
3 requesting a quality of life measure via a voice prompt played on the
4 remote client to the individual patient.

1 17. A method according to Claim 16, the operation of requesting a
2 quality of life measure further comprising:
3 storing a written script comprising a plurality of quality of life measure
4 requests within the remote client;
5 retrieving each quality of life request from the stored written script with
6 each such retrieved quality of life measure request comprising one such voice
7 prompt; and
8 synthesizing speech for playback from the retrieved quality of life request.

1 18. A method according to Claim 16, the operation of requesting a
2 quality of life measure further comprising:
3 storing pre-recorded speech comprising a plurality of quality of life
4 measure requests within the remote client;
5 retrieving each quality of life request from the stored pre-recorded speech
6 with each such retrieved quality of life measure request comprising one such
7 voice prompt; and
8 playing the pre-recorded speech from the retrieved quality of life request.

1 19. A method according to Claim 12, the operation of processing voice
2 feedback further comprising:
3 recognizing individual words in the spoken voice feedback; and
4 translating the individual spoken words into written individual words.

1 20. A method according to Claim 19, further comprising:
2 storing a voice grammar within the remote client, the voice grammar
3 comprising a plurality of speech phrases expressed in a natural language, each
4 speech phrase corresponding to a normalized quality of life measure;
5 parsing the written individual words into tokens; and
6 performing a lexical analysis of the parsed tokens in accordance with the
7 voice grammar to identify one such normalized quality of life measure.

1 21. A method according to Claim 19, further comprising:
2 storing the written individual words as a vocabulary within the remote
3 client; and
4 performing a lookup of the written individual words from the vocabulary
5 stored within the remote client.

1 22. A method according to Claim 12, wherein the remote client
2 comprises at least one of a personal computer, an audio interface, and a telephony
3 instrument.

1 23. A computer-readable storage medium holding code for providing
2 normalized voice feedback from an individual patient in an automated collection
3 and analysis patient care system, comprising:
4 periodically receiving a set of collected measures retrieved on a
5 substantially regular basis from a medical device having a sensor for monitoring
6 at least one physiological measure of an individual patient, the collected measures
7 set comprising individual measures which each relate to patient information
8 recorded by the medical device;
9 storing the collected measures set into a patient care record for the
10 individual patient within a database organized to store one or more patient care
11 records which each comprise a plurality of the collected measures sets;
12 processing voice feedback into a set of quality of life measures which each
13 relate to patient self-assessment indicators, the voice feedback having been
14 spoken by the individual patient into a remote client substantially
15 contemporaneous to the collection of an identifiable device measures set;
16 receiving the identified collected device measures set and the quality of
17 life measures set;
18 storing the identified collected device measures set and the quality of life
19 measures set into the patient care record for the individual patient within the
20 database; and
21 analyzing the identified collected device measures set, the quality of life
22 measures set, and one or more of the collected device measures sets in the patient
23 care record for the individual patient relative to one or more other collected
24 device measures sets stored in the database to determine a patient status indicator.

1 24. A storage medium according to Claim 23, further comprising:
2 repeatedly receiving one or more collected measures sets which are each
3 recorded by a sensor which monitors at least one physiological measure of the
4 individual patient, each such sensor monitoring a site within the individual patient
5 unique from the site monitored by any other such sensor;

6 storing each collected measures set organized by specific site into the
7 patient care record for the individual patient within the database; and
8 analyzing one or more of the site specific collected measures sets in the
9 patient care record for each site within the individual patient relative to one or
10 more other site specific collected measures sets stored in the database to
11 determine a patient status indicator.

1 25. A storage medium according to Claim 24, wherein the one or more
2 site specific collected measures sets and the one or more other site specific
3 collected measures sets both store measures collected from the same relative site.

1 26. A storage medium according to Claim 24, wherein the one or more
2 site specific collected measures sets and the one or more other site specific
3 collected measures sets both store measures collected from a different site.

1 27. A storage medium according to Claim 23, the operation of
2 processing voice feedback further comprising:
3 requesting a quality of life measure via a voice prompt played on the
4 remote client to the individual patient.

1 28. A storage medium according to Claim 27, the operation of
2 requesting a quality of life measure further comprising:
3 storing a written script comprising a plurality of quality of life measure
4 requests within the remote client;
5 retrieving each quality of life request from the stored written script with
6 each such retrieved quality of life measure request comprising one such voice
7 prompt; and
8 synthesizing speech for playback from the retrieved quality of life request.

1 29. A storage medium according to Claim 27, the operation of
2 requesting a quality of life measure further comprising:

3 storing pre-recorded speech comprising a plurality of quality of life
4 measure requests within the remote client;
5 retrieving each quality of life request from the stored pre-recorded speech
6 with each such retrieved quality of life measure request comprising one such
7 voice prompt; and
8 playing the pre-recorded speech from the retrieved quality of life request.

1 30. A storage medium according to Claim 23, the operation of
2 processing voice feedback further comprising:
3 recognizing individual words in the spoken voice feedback; and
4 translating the individual spoken words into written individual words.

1 31. A storage medium according to Claim 30, further comprising:
2 storing a voice grammar within the remote client, the voice grammar
3 comprising a plurality of speech phrases expressed in a natural language, each
4 speech phrase corresponding to a normalized quality of life measure;
5 parsing the written individual words into tokens; and
6 performing a lexical analysis of the parsed tokens in accordance with the
7 voice grammar to identify one such normalized quality of life measure.

1 32. A storage medium according to Claim 30, further comprising:
2 storing the written individual words as a vocabulary within the remote
3 client; and
4 performing a lookup of the written individual words from the vocabulary
5 stored within the remote client.




FIGURE 1




FIGURE 2




FIGURE 3

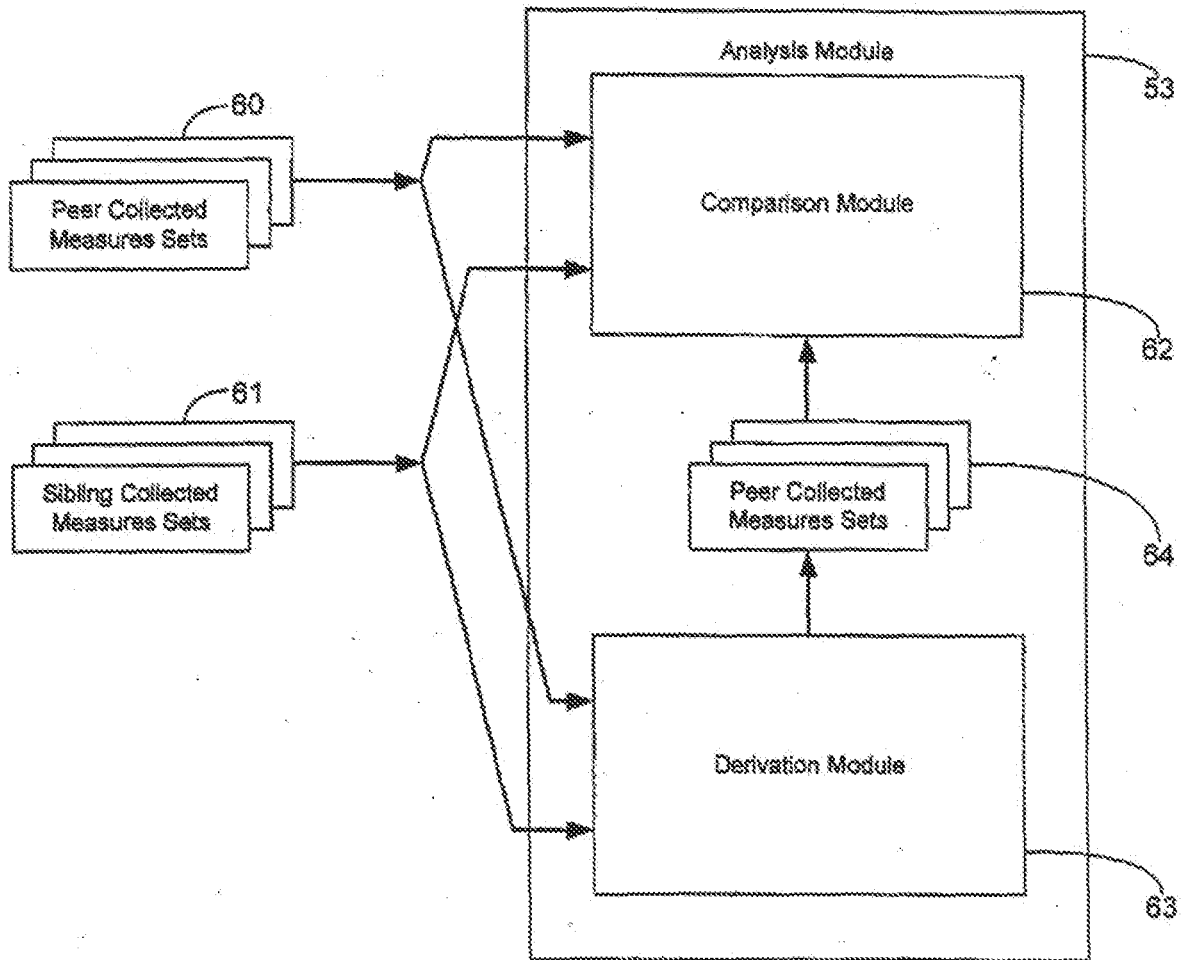


FIGURE 4




FIGURE 5

Patient 1

Set 0				Set n-2	Set n-1	Set n
X_0	*	*	*	X_{n-2}	X_{n-1}	X_n
Y_0	*	*	*	Y_{n-2}	Y_{n-1}	Y_n
Z_0	*	*	*	Z_{n-2}	Z_{n-1}	Z_n

time →

Patient 2

Set 0				Set n-2	Set n-1	Set n
X_0'	*	*	*	X_{n-2}'	X_{n-1}'	X_n'
Y_0'	*	*	*	Y_{n-2}'	Y_{n-1}'	Y_n'
Z_0'	*	*	*	Z_{n-2}'	Z_{n-1}'	Z_n'

time →

Patient 3

Set 0				Set n-2	Set n-1	Set n
X_0''	*	*	*	X_{n-2}''	X_{n-1}''	X_n''
Y_0''	*	*	*	Y_{n-2}''	Y_{n-1}''	Y_n''
Z_0''	*	*	*	Z_{n-2}''	Z_{n-1}''	Z_n''

time →

FIGURE 6

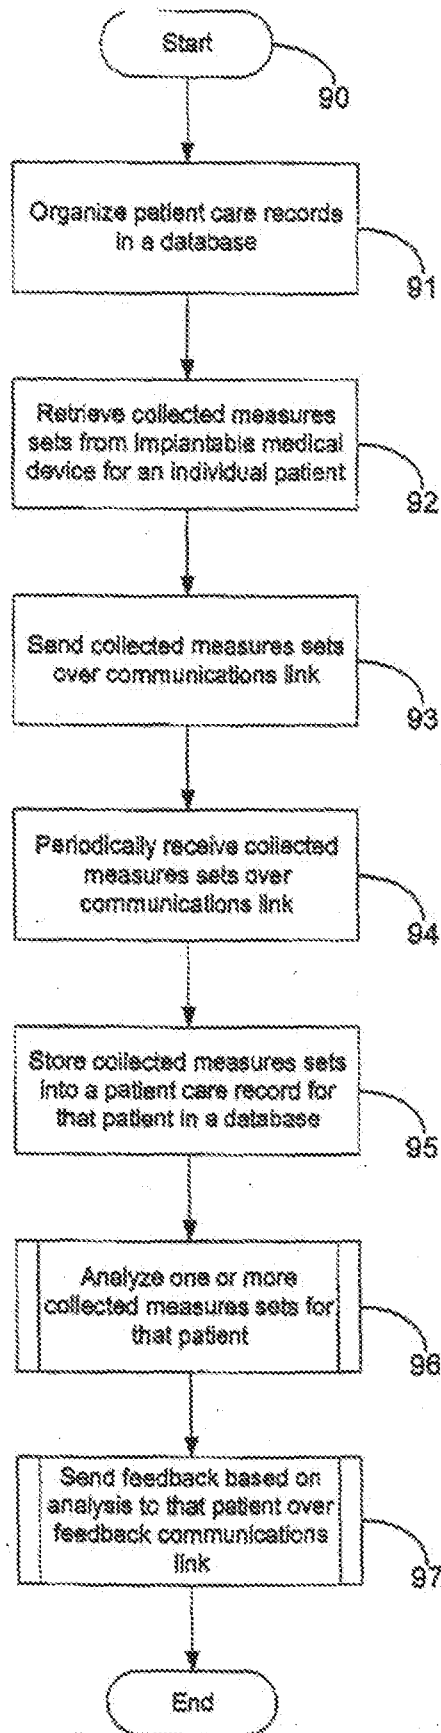


FIGURE 7




FIGURE 8

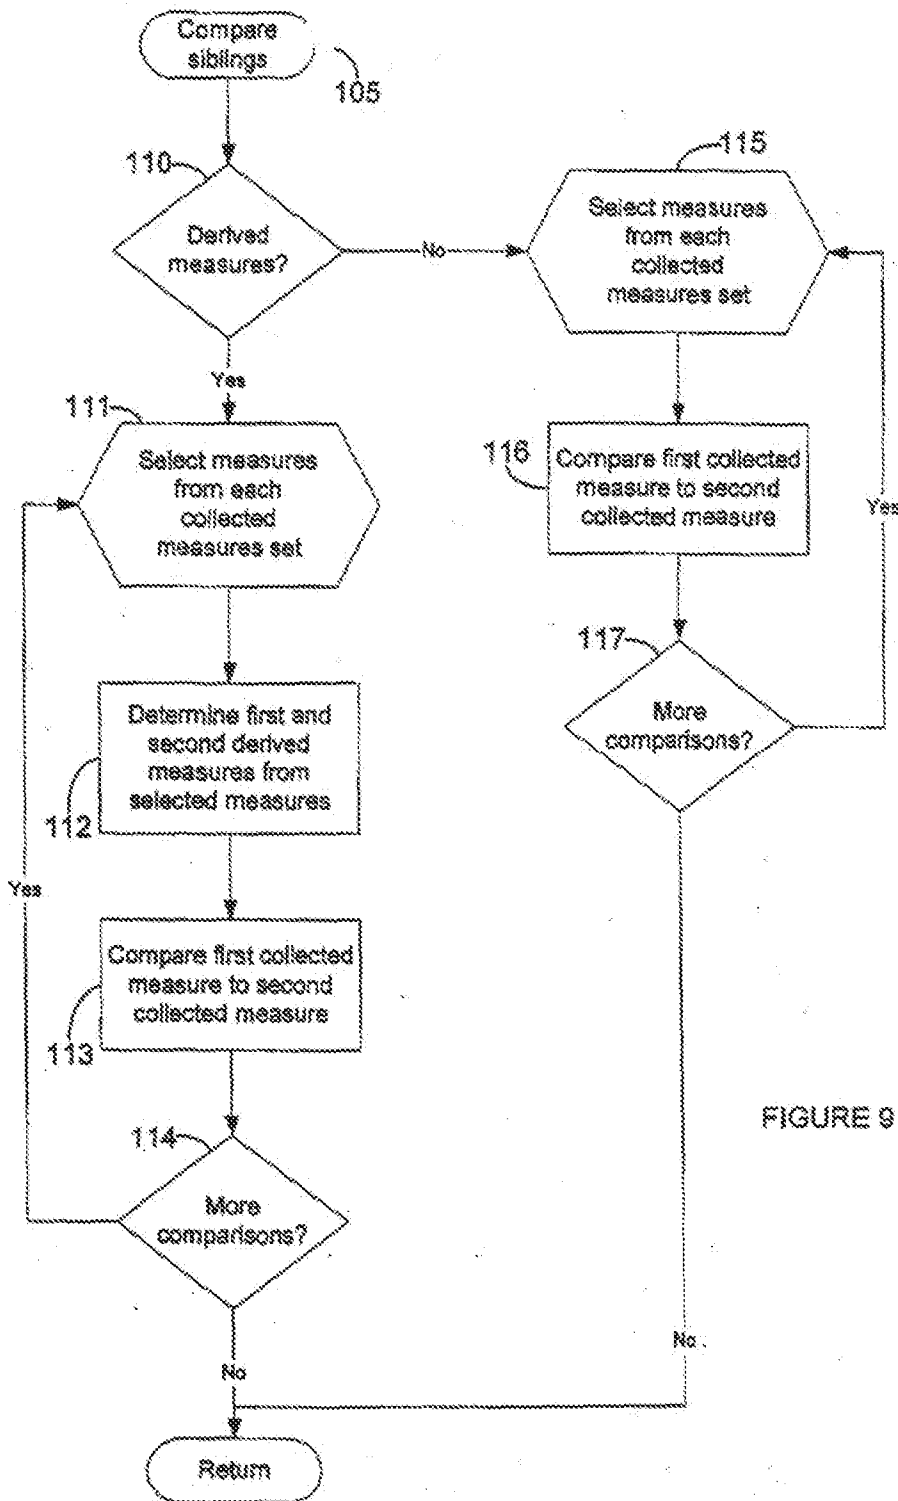


FIGURE 9




FIGURE 10A




FIGURE 10B.




FIGURE 11




FIGURE 12

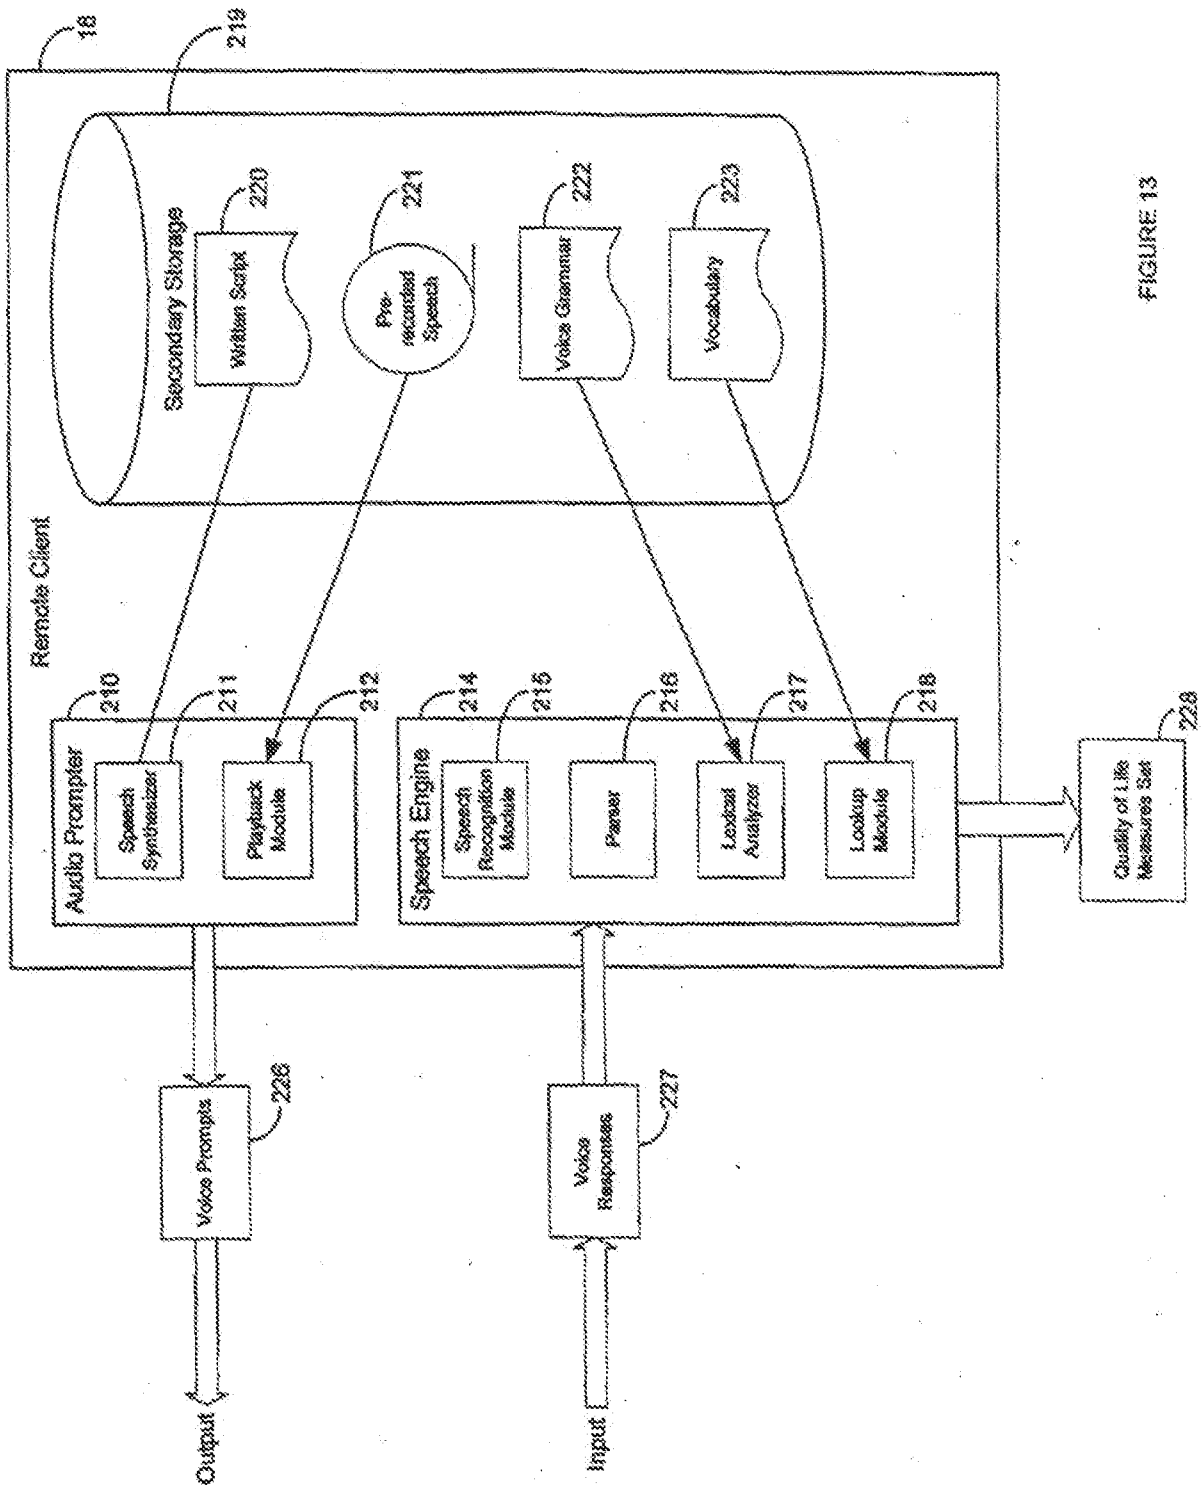


FIGURE 13




FIGURE 14




FIGURE 16

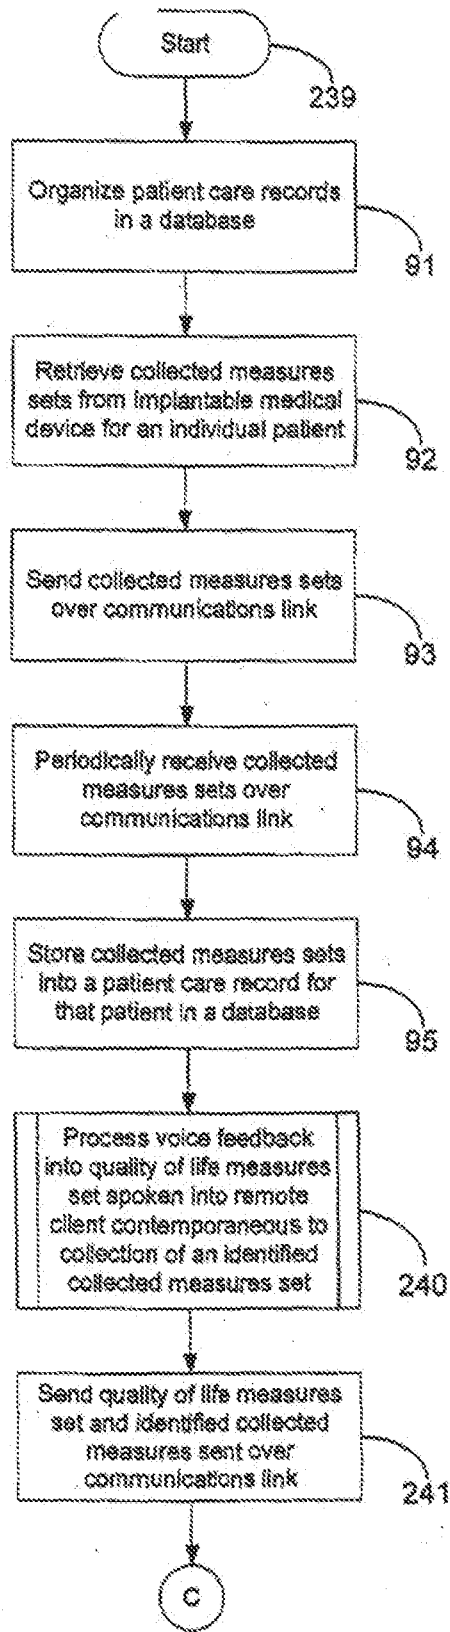


FIGURE 16A




FIGURE 18B