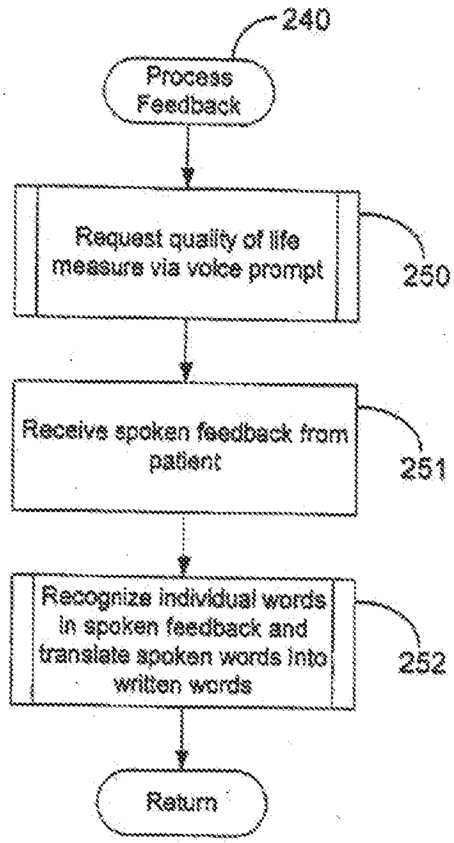
240
Process Feedback
Request quality of life measure via voice prompt
250
Receive spoken feedback from patient
251
Recognize individual words in spoken feedback and translate spoken words into written words
252
Return

FIGURE 17

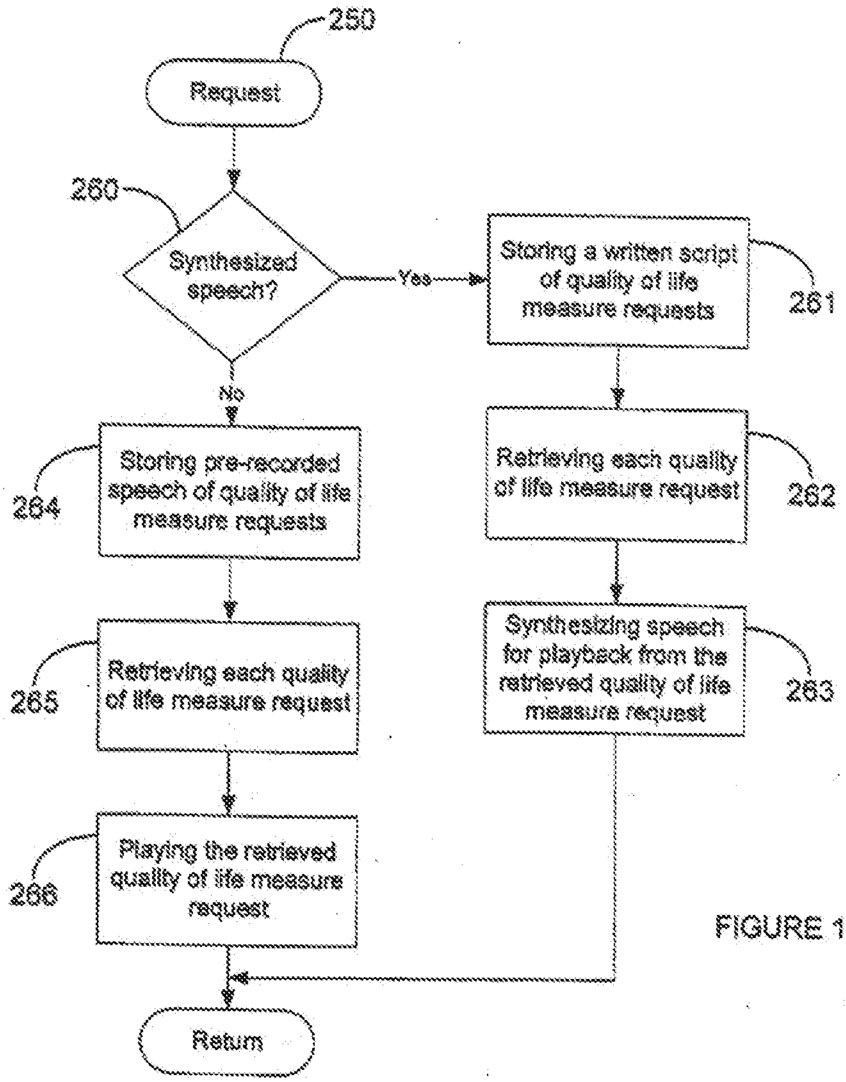
250
Request
260
Synthesized speech?
Yes
No
Storing a written script of quality of life measure requests
261
Retrieving each quality of life measure request
262
Synthesizing speech for playback from the retrieved quality of life measure request
263
Storing pre-recorded speech of quality of life measure requests
264
Retrieving each quality of life measure request
265
Playing the retrieved quality of life measure request
266
Return

FIGURE 18

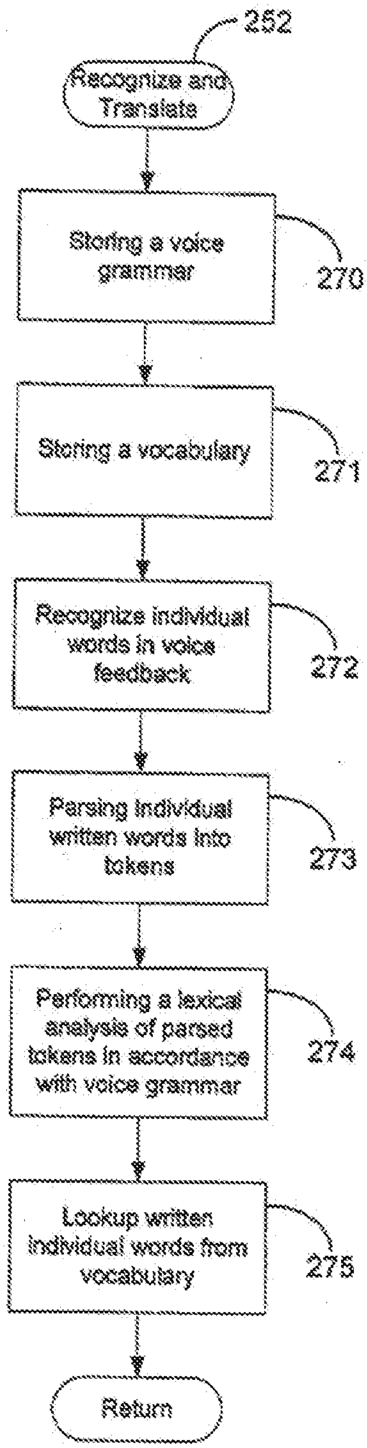
252
Recognize and Translate
Storing a voice grammar
270
Storing a vocabulary
271
Recognize individual words in voice feedback
272
Parsing individual written words into tokens
273
Performing a lexical analysis of parsed tokens in accordance with voice grammar
274
Lookup written individual words from vocabulary
275
Return

FIGURE 19

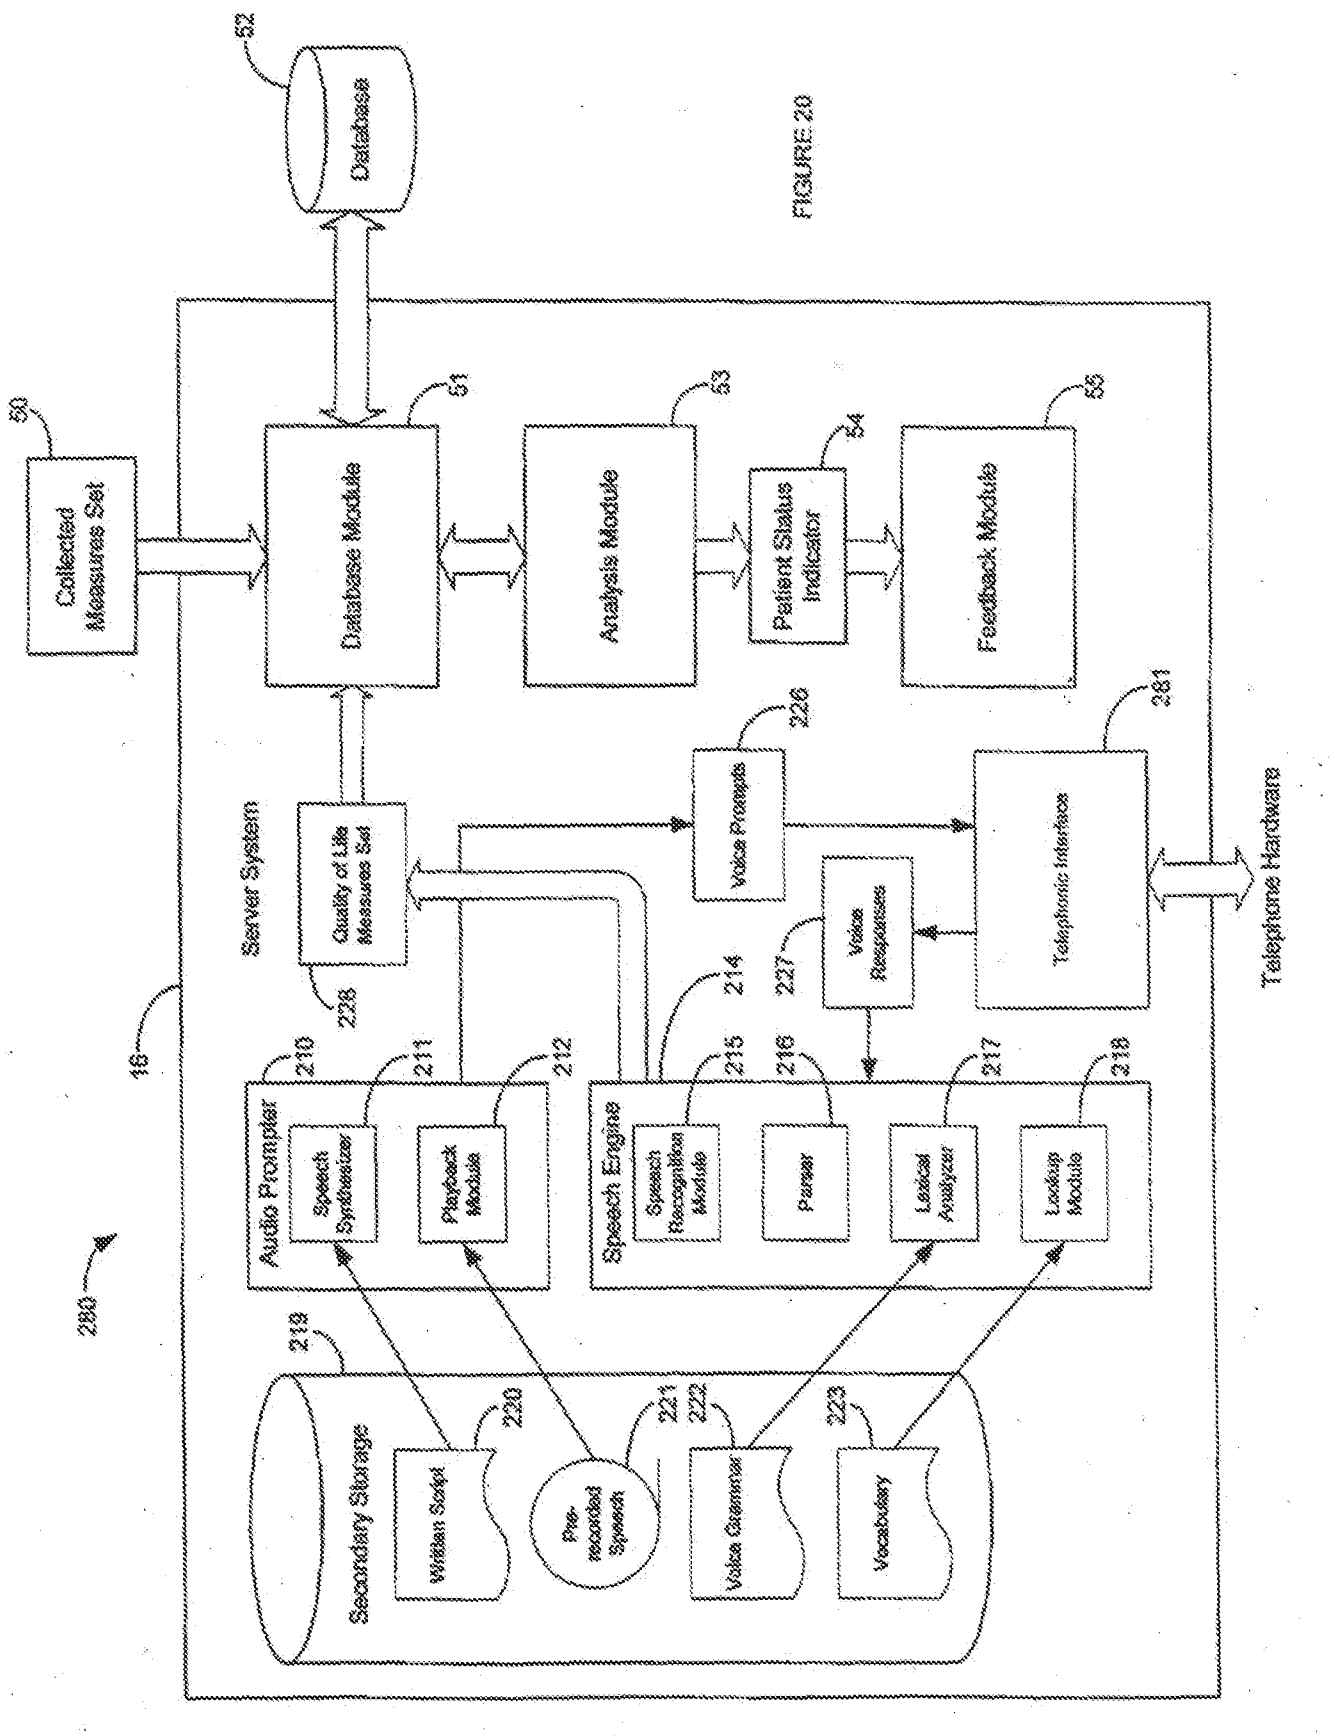
52
Database
50
Collected Measures Set
51
Database Module
53
Analysis Module
54
Patient Status Indicator
55
Feedback Module
16
Server System
228
Quality of Life Measures Set
226
Voice Prompts
281
Telephonic Interface
227
Voice Responses
Telephone Hardware
280
210
Audio Prompter
211
Speech Synthesizer
212
Playback Module
214
Speech Engine
215
Speech Recognition Module
216
Parser
217
Lexical Analyzer
218
Lookup Module
219
Secondary Storage
220
Written Script
221
Pre-recorded Speech
222
Voice Grammar
223
Vocabulary

FIGURE 20

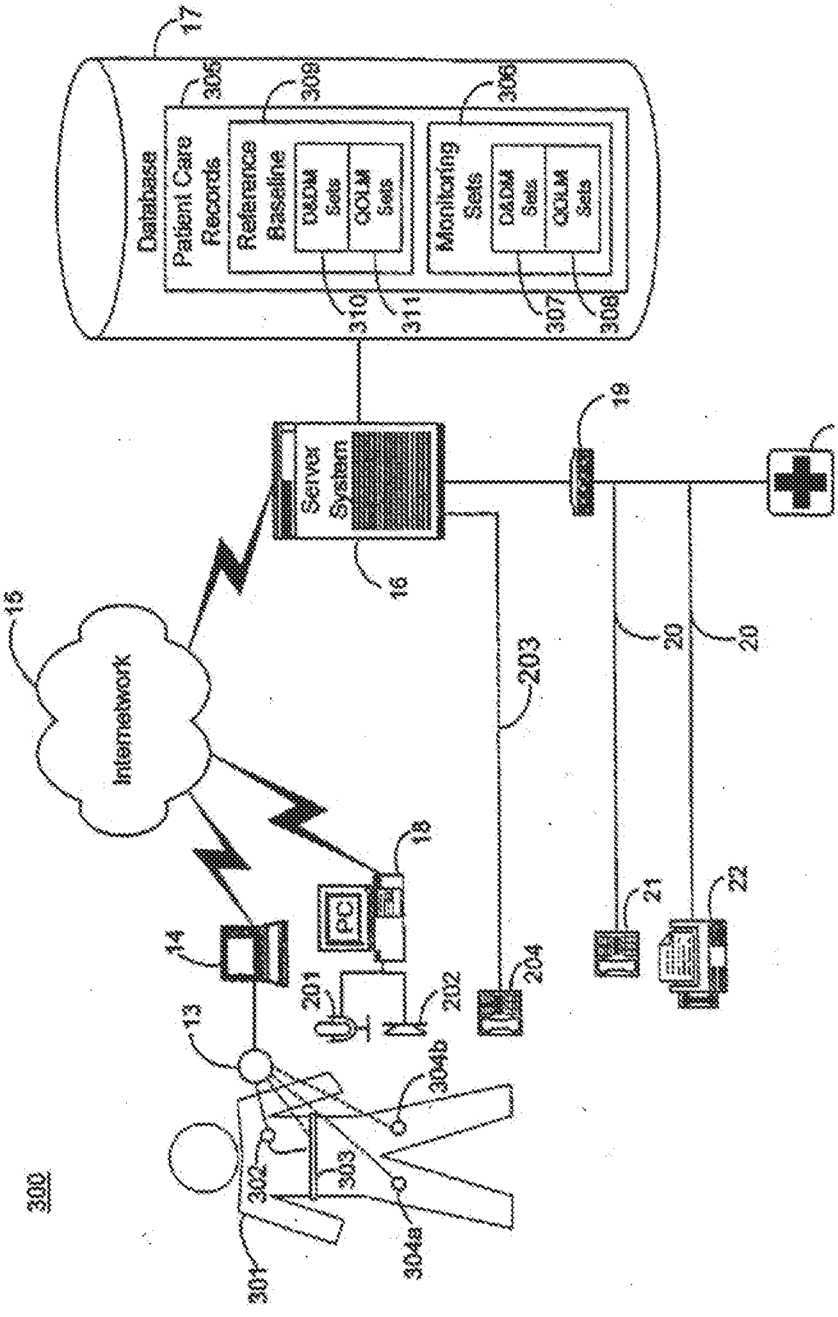
300
301
302
303
304a
304b
13
14
201
202
PC
18
Internetwork
15
Server System
16
17
Database
Patient Care Records
305
Reference Baseline
309
D&DM Sets
QOLM Sets
310
311
Monitoring Sets
306
D&DM Sets
QOLM Sets
307
308
19
312
20
20
203
204
21
22

FIGURE 21

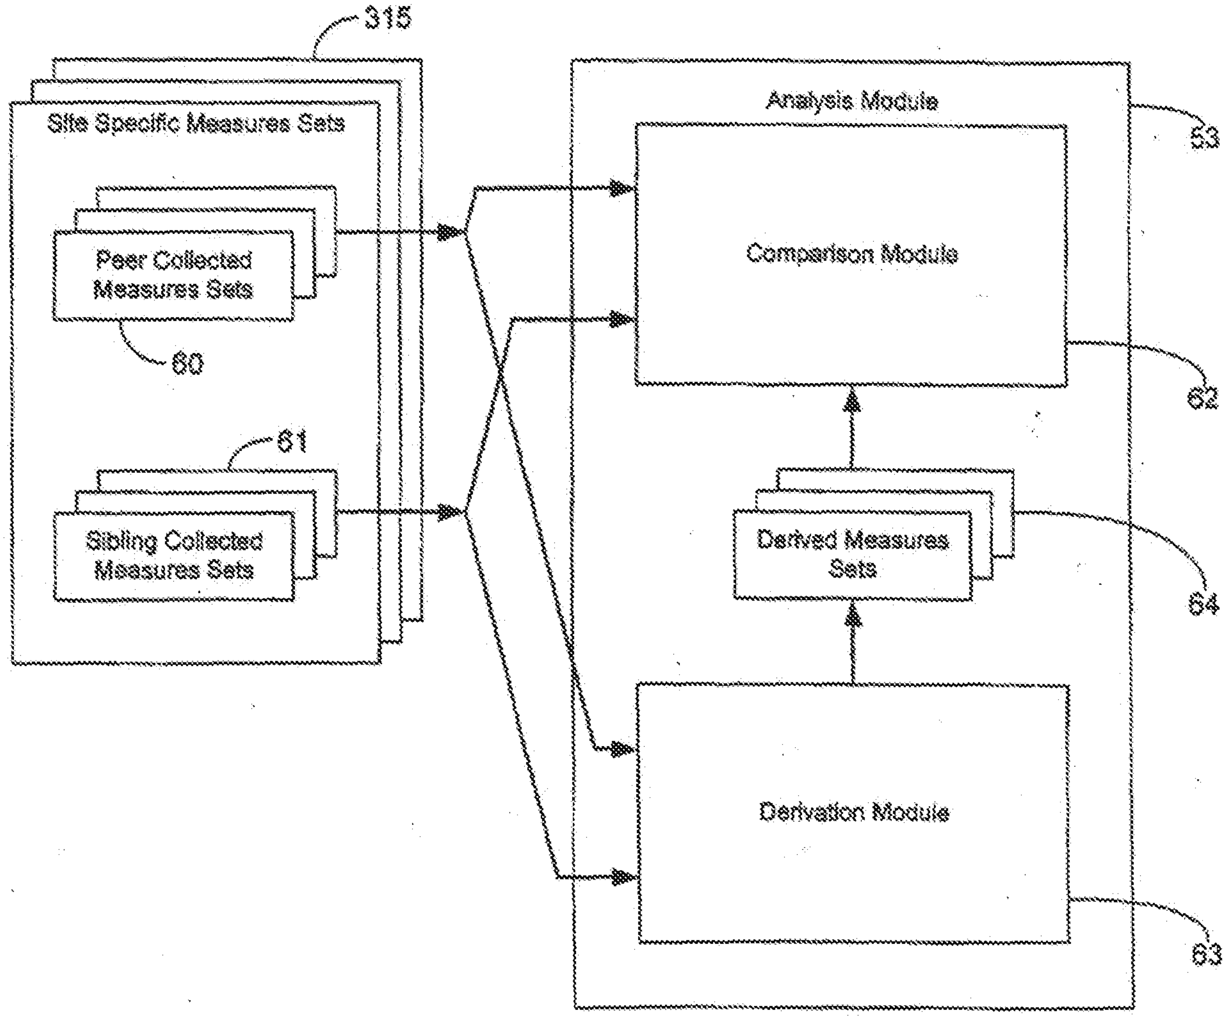
315
Site Specific Measures Sets
Peer Collected Measures Sets
60
61
Sibling Collected Measures Sets
Analysis Module
53
Comparison Module
62
Derived Measures Sets
64
Derivation Module
63

FIGURE 22

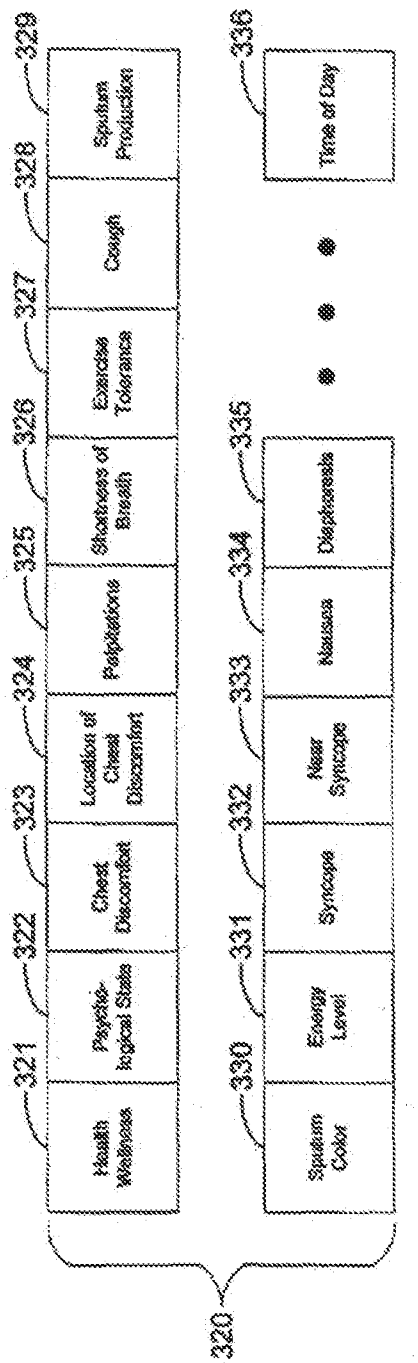
320
321
Health Wellness
322
Psycho-logical State
323
Chest Discomfort
324
Location of Chest Discomfort
325
Palpitations
326
Shortness of Breath
327
Exercise Tolerance
328
Cough
329
Sputum Production
330
Sputum Color
331
Energy Level
332
Syncope
333
Near Syncope
334
Nausea
335
Diaphoresis
336
Time of Day

FIGURE 23

Patient 1

Set 0				Set n-2	Set n-1	Set n
			Site A			
X_{0_A}	•	•	•	X_{n-2_A}	X_{n-1_A}	X_{n_A}
Y_{0_A}	•	•	•	Y_{n-2_A}	Y_{n-1_A}	Y_{n_A}
Z_{0_A}	•	•	•	Z_{n-2_A}	Z_{n-1_A}	Z_{n_A}
			Site B			
X_{0_B}	•	•	•	X_{n-2_B}	X_{n-1_B}	X_{n_B}
Y_{0_B}	•	•	•	Y_{n-2_B}	Y_{n-1_B}	Y_{n_B}
Z_{0_B}	•	•	•	Z_{n-2_B}	Z_{n-1_B}	Z_{n_B}
			Site C			
X_{0_C}	•	•	•	X_{n-2_C}	X_{n-1_C}	X_{n_C}
Y_{0_C}	•	•	•	Y_{n-2_C}	Y_{n-1_C}	Y_{n_C}
Z_{0_C}	•	•	•	Z_{n-2_C}	Z_{n-1_C}	Z_{n_C}

time →

Patient 2

Set 0				Set n-2	Set n-1	Set n
			Site A			
$X_{0'_A}$	•	•	•	$X_{n-2'_A}$	$X_{n-1'_A}$	$X_{n'_A}$
$Y_{0'_A}$	•	•	•	$Y_{n-2'_A}$	$Y_{n-1'_A}$	$Y_{n'_A}$
$Z_{0'_A}$	•	•	•	$Z_{n-2'_A}$	$Z_{n-1'_A}$	$Z_{n'_A}$
			Site B			
$X_{0'_B}$	•	•	•	$X_{n-2'_B}$	$X_{n-1'_B}$	$X_{n'_B}$
$Y_{0'_B}$	•	•	•	$Y_{n-2'_B}$	$Y_{n-1'_B}$	$Y_{n'_B}$
$Z_{0'_B}$	•	•	•	$Z_{n-2'_B}$	$Z_{n-1'_B}$	$Z_{n'_B}$

time →

Patient 3

Set 0				Set n-2	Set n-1	Set n
			Site A			
$X_{0''_A}$	•	•	•	$X_{n-2''_A}$	$X_{n-1''_A}$	$X_{n''_A}$
$Y_{0''_A}$	•	•	•	$Y_{n-2''_A}$	$Y_{n-1''_A}$	$Y_{n''_A}$
$Z_{0''_A}$	•	•	•	$Z_{n-2''_A}$	$Z_{n-1''_A}$	$Z_{n''_A}$
			Site D			
$X_{0''_D}$	•	•	•	$X_{n-2''_D}$	$X_{n-1''_D}$	$X_{n''_D}$
$Y_{0''_D}$	•	•	•	$Y_{n-2''_D}$	$Y_{n-1''_D}$	$Y_{n''_D}$
$Z_{0''_D}$	•	•	•	$Z_{n-2''_D}$	$Z_{n-1''_D}$	$Z_{n''_D}$

time →

FIGURE 24

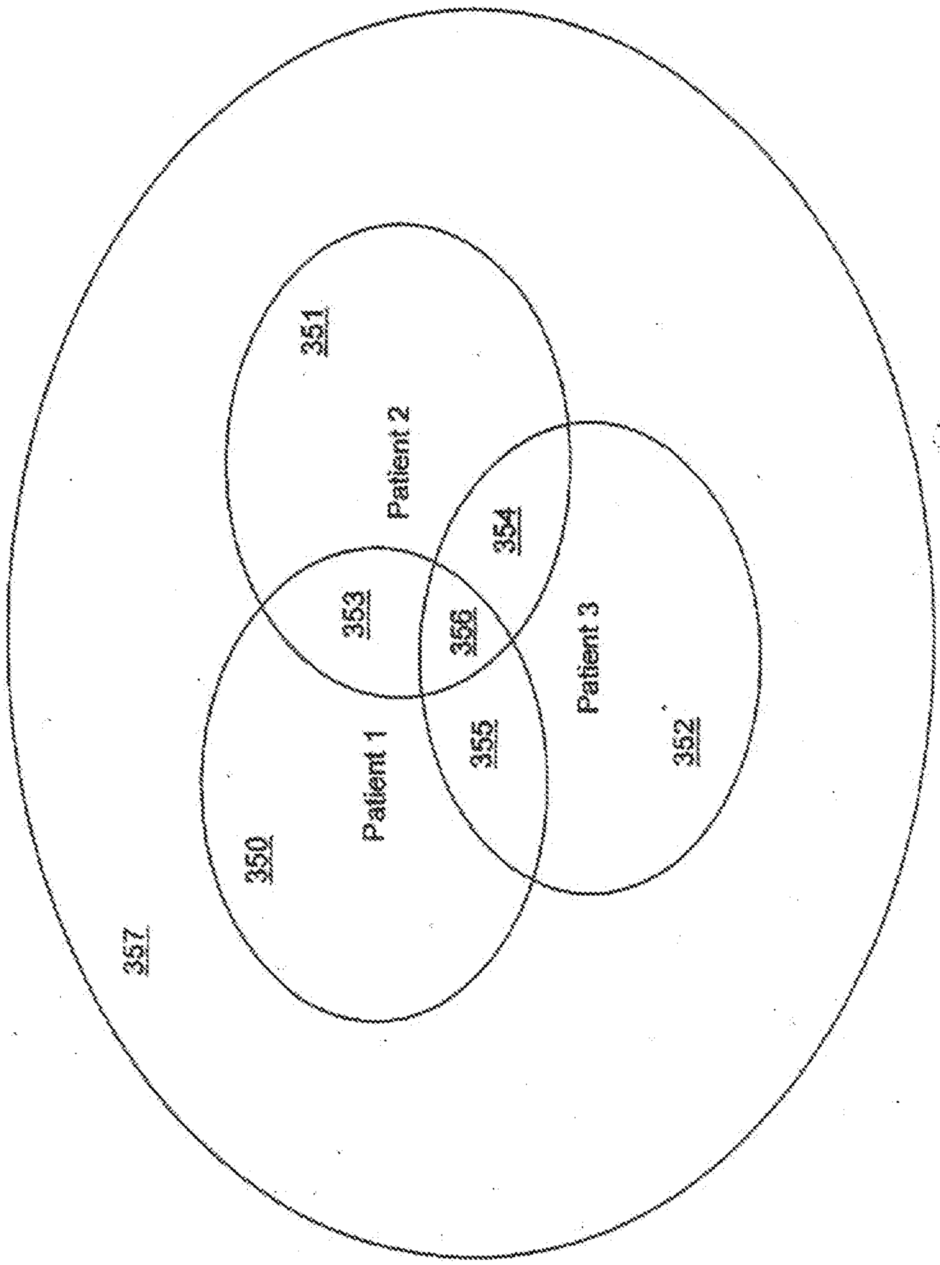

FIGURE 25

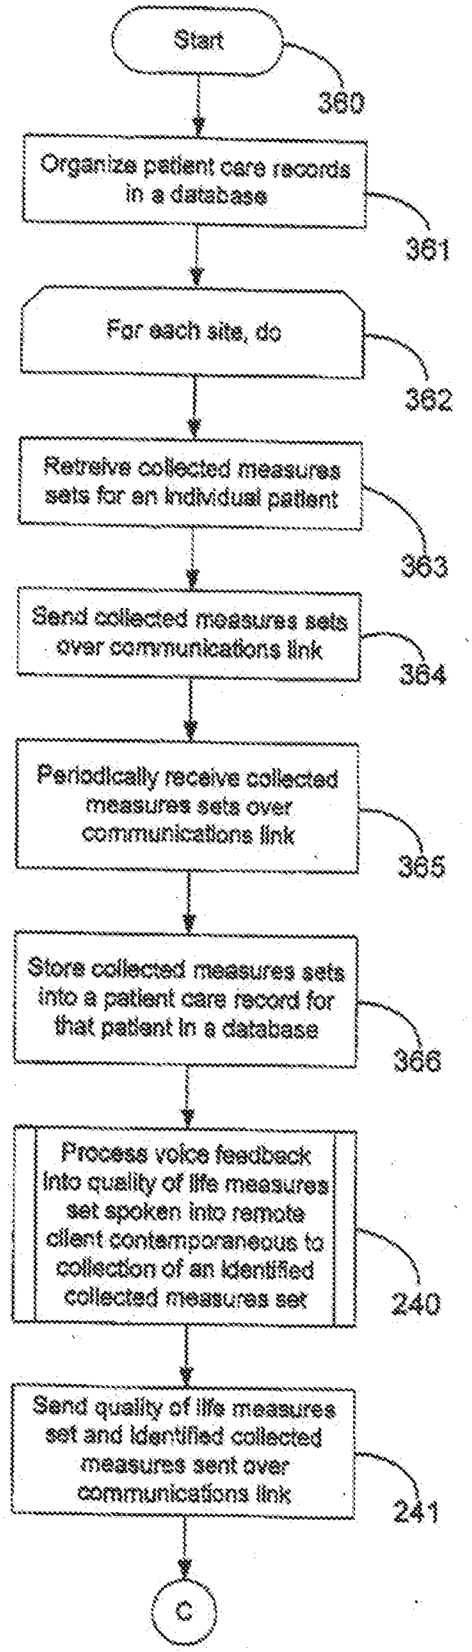
Start
360
Organize patient care records in a database
361
For each site, do
362
Retreive collected measures sets for an individual patient
363
Send collected measures sets over communications link
364
Periodically receive collected measures sets over communications link
365
Store collected measures sets into a patient care record for that patient in a database
366
Process voice feedback into quality of life measures set spoken into remote client contemporaneous to collection of an identified collected measures set
240
Send quality of life measures set and identified collected measures sent over communications link
241
C

FIGURE 26A

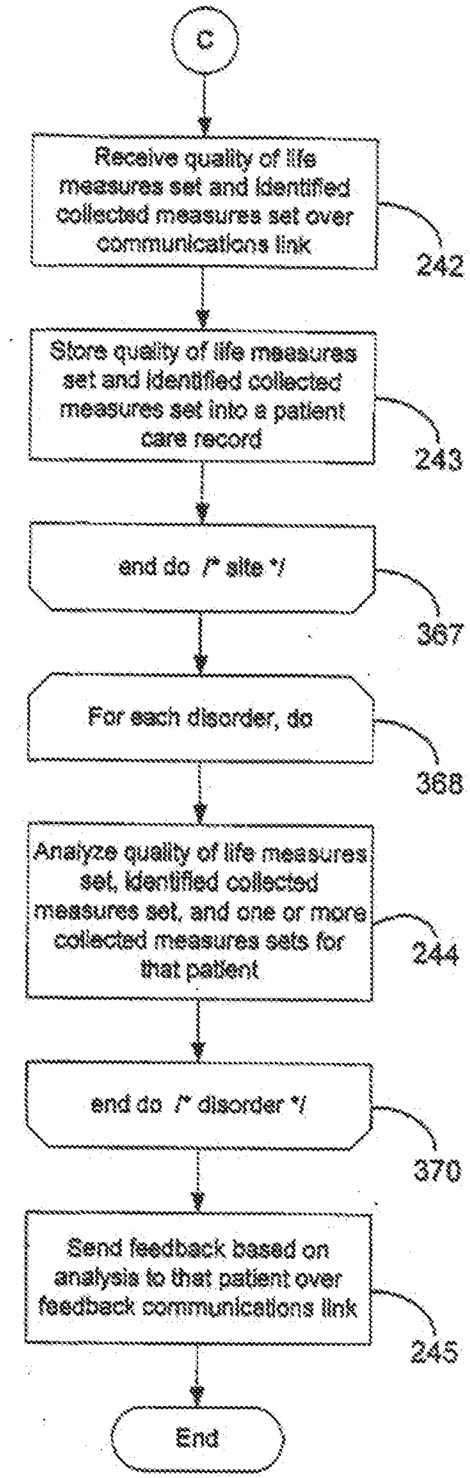
C
Receive quality of life measures set and identified collected measures set over communications link
242
Store quality of life measures set and identified collected measures set into a patient care record
243
end do /* site */
367
For each disorder, do
368
Analyze quality of life measures set, identified collected measures set, and one or more collected measures sets for that patient
244
end do /* disorder */
370
Send feedback based on analysis to that patient over feedback communications link
245
End

FIGURE 26B

(19) Canadian Intellectual Property Office
An Agency of Industry Canada

Office de la Propri,t, intellectuelle du Canada
Un organisme d'industrie Canada

(11) CA 2 387 039 (13) A1

(40) 19.04.2001
(43) 19.04.2001

(12)

(21) 2 387 039

(22) 13.10.2000

(51) Int. Cl. 7: G06F 15/16, G06F 15/173, G06F 15/177

(85) 08.04.2002

(86) PCT/US00/28514

(87) WO01/027783

(30) 60/159,186 US 13.10.1999

(71)
SIRENIC, INC.,
2109 Landings Drive, MOUNTAIN VIEW, XX (US).

(72)
CALLENDER, JOHN LENNON (US).
SHEARD, NICOLAS (GB).

(74)
SMART & BIGGAR

(54) SYSTEME ET PROCEDES PERMETTANT D'ACCEDER A DES INFORMATIONS SUR INTERNET PAR UTILISATION DE DISPOSITIFS INTERNET

(54) SYSTEM AND METHODS FOR ACCESSING INTERNET INFORMATION USING INTERNET APPLIANCES

(57)
A system and methods for accessing summary or pre-selected items of Internet information using Internet appliances (20-25) is provided. Users access Internet information through an Internet appliance interface optimized for the network bandwidth and display capabilities (22) of the Internet appliance. The Internet information accessed is the information that is determined to be relevant and important to the user. The information is summarized for, and presented to, the user at one or more pre-determined levels of detail.

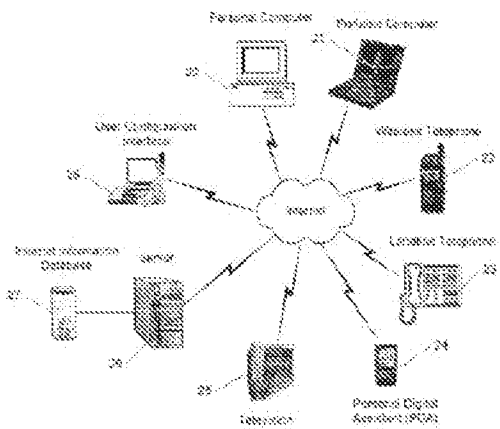

Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2387039 A1 2001/04/19

(21) 2 387 039

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(86) Date de dépôt PCT/PCT Filing Date: 2000/10/13
(87) Date publication PCT/PCT Publication Date: 2001/04/19
(85) Entrée phase nationale/National Entry: 2002/04/05
(86) N° demande PCT/PCT Application No.: US 2000/028514
(87) N° publication PCT/PCT Publication No.: 2001/027763
(30) Priorité/Priority: 1999/10/13 (60/159,186) US
(51) Cl.Int.⁷/Int.Cl.⁷ G06F 15/16, G06F 15/177, G06F 15/173

(71) Demandeur/Applicant:
SIRENIC, INC., US
(72) Inventeurs/Inventors:
SHEARD, NICOLAS, GB;
CALLENDER, JOHN LENNON, US
(74) Agent: SMART & BIGGAR

(54) Titre : SYSTEME ET PROCEDES PERMETTANT D'ACCEDER A DES INFORMATIONS SUR INTERNET PAR UTILISATION DE DISPOSITIFS INTERNET
(54) Title: SYSTEM AND METHODS FOR ACCESSING INTERNET INFORMATION USING INTERNET APPLIANCES

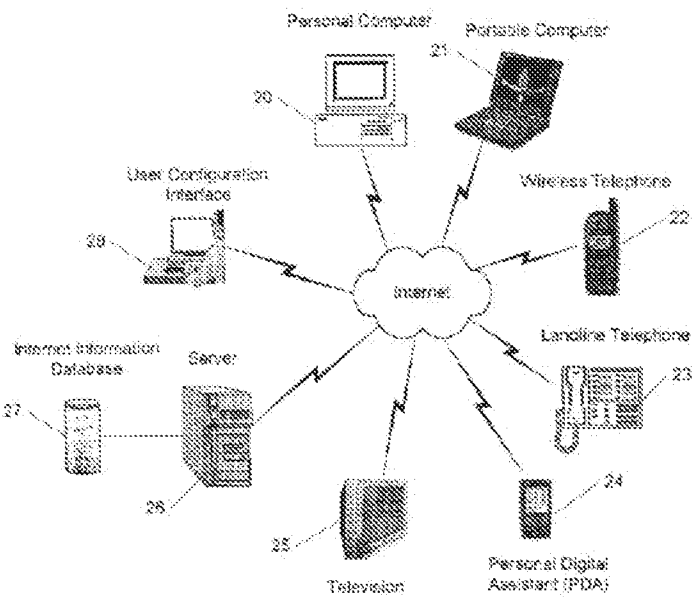

(57) Abrégé/Abstract:
A system and methods for accessing summary or pre-selected items of Internet information using internet appliances (20-25) is provided. Users access internet information through an internet appliance interface optimized for the network bandwidth and display capabilities (22) of the Internet appliance. The internet information accessed is the information that is determined to be relevant and important to the user. The information is summarized for, and presented to, the user at one or more pre-determined levels of detail.

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
19 April 2001 (19.04.2001)

PCT

(10) International Publication Number
WO 01/27783 A1

(51) International Patent Classification[7]: G06F 15/16, 15/173, 15/177

(21) International Application Number: PCT/US00/28514

(22) International Filing Date: 13 October 2000 (13.10.2000)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/159,186 13 October 1999 (13.10.1999) US

(71) Applicant: SIRENIC, INC. [US/US]; 2350 W. El Camino Real, Suite 210, Mountain View, CA 94040-1456 (US).

(72) Inventors: SHEARD, Nicolas; 5 Hinton Wood Avenue, Highcliffe, Christchurch, Dorset BH23 5AB (GB). CALLENDER, John, Lennon; 839 Rockwood Drive, San Jose, CA 95139 (US).

(74) Agents: PISANO, Nicola, A. et al.; Fish & Neave, 1251 Avenue of the Americas, New York, NY 10020 (US).

(81) Designated States (*national*): AU, CA, JP.

(84) Designated States (*regional*): European patent (AT, BE, CH, CY, DE, DK, ES, FI, FR, GB, GR, IE, IT, LU, MC, NL, PT, SE).

Published:
— *With international search report.*

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

(54) Title: SYSTEM AND METHODS FOR ACCESSING INTERNET INFORMATION USING INTERNET APPLIANCES

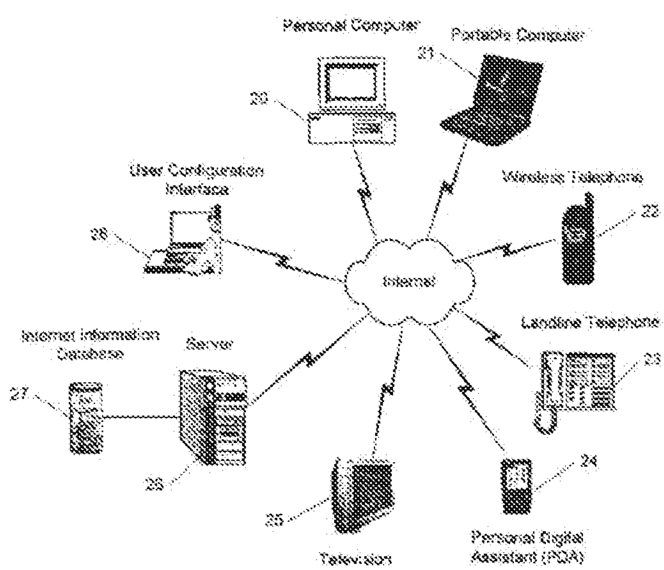

(57) Abstract: A system and methods for accessing summary or pre-selected items of Internet information using Internet appliances (20-25) is provided. Users access Internet information through an Internet appliance interface optimized for the network bandwidth and display capabilities (22) of the Internet appliance. The Internet information accessed is the information that is determined to be relevant and important to the user. The information is summarized for, and presented to, the user at one or more pre-determined levels of detail.

WO 01/27783 A1

SYSTEM AND METHODS FOR ACCESSING INTERNET INFORMATION USING INTERNET APPLIANCES

Field Of The Invention

This invention relates generally to a system and methods for providing access to Internet information from Internet appliances. More specifically, the present invention provides a system and methods for accessing summary or pre-selected items of Internet from Internet appliances according to users' preferences.

Background Of The Invention

The Internet and the World Wide Web (hereinafter "the web") have revolutionized the ways in which information is disseminated and shared. A wide variety of information can be simultaneously accessed by multiple users through a new category of documents designed to easily represent content for display and transmission over the Internet. These new documents, often referred to as *electronic documents* or *web pages*, are increasingly replacing their traditional paper counterparts as the medium through which business is carried out.

A web page is a multimedia composition that is displayed to the user on a "web browser window" by "web browser software". Under the control of a user, the web browser software establishes a connection over the Internet between the user's computer, and a "web server". This connection is used to download data representing a "web page" from the web server to the user's computer. Web pages may contain text, audio, graphics, imagery, and video content, as well as nearly any other type of content that may be experienced through use of a computer or other electronic devices. Additionally, web pages may be interactive, and may contain user selectable links that cause other web pages to be displayed, forms that may be used to send information from the user to the web server, interactive executable code, or other elements through which the user may interact with web pages. A group of one or more interconnected and closely related web pages, such as all the web pages containing information about a single company, located on one or more web servers, is referred to as a "web site".

At present, information displayed on web pages in the Internet can be accessed by users from various "Internet appliances", which are electronic devices configured with Internet access systems. Internet appliances include, but are not limited to, microprocessor based devices such as personal and portable computers, personal digital assistants, electronic organizers and toys, as well as consumer media delivery devices such as land line and wireless mobile telephones, television, and radio. The Internet access systems allow users to access information such as their e-mail, news, stock quotes, among others, from different Internet appliances. The information is accessed through an interface, which can be either a *visual interface* or a *voice interface*, depending on the Internet appliance being used to access the information.

A visual interface may be used in Internet appliances that have a visual display, such as personal and portable computers, personal digital assistants (PDA), electronic organizers and toys, as well as consumer media delivery devices such as wireless mobile telephones and television. In appliances that have a small visual display such as PDAs, electronic organizers, and wireless telephones, access to Internet information is provided via *microbrowser* software, which consists of a simpler version of a web browser, with reduced graphic capabilities. Users select an information content provider in the Internet from a pull-down menu displayed in the microbrowser window in the wireless telephone. Such content providers include Yahoo! from Santa Clara, CA, Excite from Redwood City, CA, America Online, Inc., from Dulles, VA, as well electronic commerce companies such as Amazon.com from Seattle, WA. The user makes a specific information request to the content provider, and the requested information is then delivered to the user by the microbrowser in the telephone display.

A voice interface may be used in Internet appliances with speech recognition technologies, such as land line and wireless mobile telephones, to access an *interactive voice response* (IVR) system. These systems enable a telephone user to access information by dialing a phone number corresponding to an IVR system server. The server contains software to issue voice prompts to the user corresponding to several information access options, and the user responds to the voice prompts by either selecting and pushing buttons on the telephone, or by using speech recognition technologies. Users can access IVR systems from any standard telephone, including those that do not have visual displays.

Examples of IVR systems include those employed by financial institutions to give users access to their accounts over the telephone, such as the system covered in U.S. Patent Number 5,825,856. Similar systems are also employed in a host of other applications, including those in the airline industry offering users the ability to check flight information over the telephone, and the system developed by Tellme Networks, Inc., of Mountain View, CA, that allows users to access information on a variety of topics, such as weather, traffic reports, restaurant reviews, stock quotes, among others.

Both visual and voice interface systems utilize a *request-wait-response* mechanism to access information that may be interesting to a user from an Internet appliance. The user makes requests on a very specific topic, waits for the information to be collected by the system from a server, and receives the information in the visual display of the Internet appliance in case a visual interface is used, or in the form of voice prompts if a voice interface is used. The request-wait-response mechanism works well on high network and display bandwidth Internet appliances such as personal computers, where users can quickly sift through information and make further requests to find exactly the information they are looking for. However, the result is less satisfactory on low network and display bandwidth Internet appliances, where the speed of information delivery is slower, presentation capabilities are less effective, and control machinery to interact and search for more relevant information is difficult to use.

As an example, consider the case of users wishing to examine their stock portfolio through a wireless telephone. Using existing systems such as the visual-based system designed for wireless telephones provided by Sprint Co. from Westwood, KS, and the voice-based system provided by Tellme Networks, Inc., of Mountain View, CA, the users may either individually request details on each company in the stock portfolio by using the telephone's limited visual or voice interface, or the users may request a list of all the stock information as a single stock portfolio. In both cases, the systems do not provide access to summary or pre-selected items of information quickly and easily. Different pieces of information are often relevant to different users, and users may have to iterate through several items of less important and less interesting information before accessing desired information.

The importance of the information may be dependent on a combination of complex parameters, for example, it may be determined by a combination of how much money users have in the stock in question, what is the percentage change in the stock price, what is the day's range, what is the volume, and when the quarterly results for the company issuing the stock are announced. Users must sift through this information using the limited interface in the telephone, which is a very time consuming and laborious process. Further, the telephone and Internet connections may be unreliable, and the user may have to wait a significant amount of time to receive the requested information.

In view of the foregoing drawbacks of accessing Internet information from Internet appliances, it would be desirable to provide a system and methods for accessing summary and pre-selected items of Internet information quickly and easily from Internet appliances according to users' preferences.

It further would be desirable to provide a system and methods for creating an Internet information database to store Internet information in categories, manage the relative importance of different items of information, and manage the relationships between different items of information.

It still further would be desirable to provide a system and methods for accessing desired information that is relevant to a particular user's interests at any given moment in time quickly within the Internet information database.

It also would be desirable to provide a system and methods for creating users' profiles to manage the information interests of each user in different information categories.

It further would be desirable to provide a system and methods for delivering summary and pre-selected items of information desired by a user in a concise and easy to comprehend format optimized for the network bandwidth and display capabilities of each individual Internet appliance.

Summary Of The Invention

In view of the foregoing, it is an object of the present invention to provide a system and methods for accessing summary and pre-selected items of Internet information quickly and easily from Internet appliances according to users' preferences.

It is another object of the present invention to provide a system and methods for creating an Internet information database to store Internet information in categories, manage the relative importance of different items of information, and manage the relationships between different items of information.

It is a further object of the present invention to provide a system and methods for accessing the desired information that is relevant to a particular user's interests at any given moment in time quickly within the Internet information database.

It is also an object of the present invention to provide a system and methods for creating users' profiles to manage the information interests of each user in different information categories.

It is a further object of the present invention to provide a system and methods for delivering summary and pre-selected items of information to a user in a concise and easy to comprehend format optimized for the network bandwidth and display capabilities of each individual Internet appliance.

These and other objects of the present invention are accomplished by providing a system and methods for accessing summary and pre-selected items of Internet information from Internet appliances quickly and easily according to a user's preferences. The user preferences are stored in a user profile that specifies the information interests of the user in a variety of information categories, such as weather, news, stock quotes, among others. The information is stored in an Internet information database designed to enable an

Internet appliance to access desired information for a given user's interests as well as other desired related information. Information determined to be desired by a user is delivered in a format optimized for network bandwidth and display capabilities of that user's Internet appliance. Such a system and methods may automatically provide more detail on information that is deemed relevant to an user, while less important and less interesting information is provided only in a summary format.

In a preferred embodiment, the system of the present invention for providing access to pre-selected items of Internet information from Internet appliances according to users' preferences involve four components: (1) an information server; (2) an information database; (3) a user configuration interface; and (4) an Internet appliance interface.

The information server executes the system software and contains four software modules: (1) an information retrieval module; (2) an information analysis module; (3) an information classification module; and (4) an information rendering module. The server also stores users' profiles that specify the information interests of each user in a variety of information categories, such as weather, news, stock quotes, among others.

The information retrieval module retrieves information from diverse Internet resources such as web servers, news feeds, and e-mail servers. The information analysis module parses the retrieved information to extract different representations at varying levels of detail. The information classification module, classifies the extracted information according to its category and priority.

The information rendering module then extracts the information determined to be desired by a user from the information database, and renders that information at one or more levels of detail and summarization dependent on that user's interest in the information, the priority of the information, and the information rendering capabilities of the user's Internet appliance. The renderings of individual items of information, which may be on a variety of categories and at different levels of detail and summarization, are combined into a summary delivered to the user in a visual format and presentation optimized for the user's Internet appliance. Optionally, the information may be transformed to facilitate presentation using a voice interface.

The information database is designed to enable an Internet appliance to access information according to a given user's interests and as well as other related information. The database records the information in categories and provides mechanisms to determine the relationships between information in different categories. These relationships define sub-categories of information and clusters of information that group related categories and sub-categories together. The clusters of information may be defined by manually programmed relationships between information categories and sub-categories, as well as by relationships formed when natural language processing of the information database can determine what information categories and sub-categories associated with the information accessed by a user are available.

The user specifies which information he or she considers relevant through a user configuration interface. The interface enables the user to explicitly specify not only the categories of information of interest, but also the relative level of interest in each information category. The information is accessed through an Internet appliance interface in a concise and easy to comprehend format optimized for network bandwidth and display capabilities of each individual Internet appliance.

Advantageously, the present invention enables a user to access Internet information quickly and easily from any Internet appliance.

In addition, the present invention enables a user to select the information accessible through the Internet appliance according to his or her personal preferences.

Brief Description Of The Drawings

The foregoing and other objects of the present invention will be apparent upon consideration of the following detailed description, taken in conjunction with the accompanying drawings, in which like reference characters refer to like parts throughout, and in which:

FIG. 1 is a schematic view of the system and the network environment in which the present invention operates;

FIG. 2 is a flow chart of the methods employed by the information server of the present invention;

FIG. 3 is a schematic view of the transactional queuing mechanisms for retrieving Internet information;

FIG. 4 is a schematic view of an illustrative information query;

FIG. 5 shows a logic equation to calculate the importance and priority of a weather report;

FIG. 6 is a schematic view of an illustrative fragment of a categorization tree for managing relationships between information items in the information database;

FIG. 7 is a schematic view of an illustrative ad hoc relationship between content nodes in a categorization tree;

FIG. 8 is a schematic view of an illustrative user profile;

FIG. 9 shows an equation to calculate the importance of a content or virtual content node in a categorization tree to an user;

FIG. 10 is a schematic view of an illustrative configured rendering definition for the lottery information category;

FIG. 11 is a flow chart for combining renderings of information categories in accordance with the principles of the present invention;

FIG. 12 is a flow chart for transforming text to facilitate presentation using a voice interface;

FIG. 13 is a schematic view of illustrative written text to voice transformations for improving the understandability of text for voice-based Internet appliance interfaces;

FIG. 14 shows an example text transcript of user access to an illustrative voice-based Internet appliance interface; and

FIGS. 15A and 15B are schematic views of illustrative visual-based interfaces for use with Internet appliances having a visual display.

Detailed Description Of The Invention

The present invention provides a system and methods for accessing concise, relevant, and important items of Internet information quickly and easily from Internet appliances according to an user's preferences. Internet appliances, as used herein, are electronic devices configured with Internet access systems. Internet appliances include, but are not limited to, microprocessor based devices such as personal and portable computers, personal digital assistants, electronic organizers and toys, as well as consumer media delivery devices such as land line and wireless mobile telephones, television, and radio.

Referring to FIG. 1, a schematic view of the system and the network environment in which the present invention operates is described. Internet appliances 20 (personal computer), 21 (portable computer), 22 (mobile telephone), 23 (land-line telephone), 24 (PDA), and 25 (television) communicate across the Internet with information server 26. Internet appliances 20-25 have an Internet appliance interface to access Internet information in a concise and easy to comprehend format optimized for the network bandwidth and display capabilities of each individual appliance.

The Internet appliance interface may be a *visual interface* or a *voice interface*, depending on the Internet appliance being used to access the information. A visual interface may be used in Internet appliances that have a visual display, such as personal computer 20 and portable computer 21, wireless mobile telephone 22, PDA 24, and television 25. A voice interface may be used in Internet appliances with speech recognition technologies, such as wireless mobile telephone 22 and land line telephone 23.

The Internet appliance interfaces access Internet information by connecting to information server 26, which executes the system software and is responsible for retrieving information from diverse Internet resources such as web servers, news feeds, and e-mail servers. Information server 26 parses the retrieved information to extract different representations at varying levels of detail and classify the extracted representations according to their category and priority.

Information server 26 also stores users' profiles that specify the information interests of each user in a variety of information categories, such as weather, news, stock quotes, among others, and extracts the information that is currently most important to the user from information database 27. Information server 26 combines the extracted information into a concise and easy to understand summary to be delivered to the user's Internet appliance.

Information database 27 enables Internet appliances 20-25 to access the most important information for a given user's interests. Information database 27 stores the information retrieved by information server 26 in categories and provides mechanisms to determine the relationships between information in different categories.

Users of Internet appliances 20-25 specify which information they consider to be important and relevant to them through user configuration interface 28. User configuration interface 28 enables users to explicitly specify not only the categories of information in which they are interested, but also their relative level of interest in each information category. User configuration interface 28 may be a user configuration web site, where a questionnaire about the user's interests can be completed. Information collected in the questionnaire is transmitted to information server 26, which translates the collected information into a user profile. User configuration interface 28 also may be a call center where a user talks to a human operator who asks questions about the user to build up the user profile.

It will be apparent to one of ordinary skill in the art that Internet appliances 20-25 are illustrative, and that alternative Internet appliances may be used with the present invention. It will further be understood that the present invention also could be used in other network settings. For example, rather than connecting through the Internet, the system and methods of the present invention could be used on a local area network. In such a configuration, the Internet appliances and server would all be connected to the same local area network.

I. Overview of Information Flow from the Internet to an Internet Appliance

Referring to FIG. 2, a flow chart of the methods employed by the information server to deliver concise, relevant, and important information from the Internet to an Internet appliance is described. At step 30, the information retrieval module of the information server reliably retrieves information from diverse Internet resources such as web servers, news feeds, and e-mail servers. The information may contain news and weather reports, stock quotes, e-mails, among others. The retrieved information is parsed at step 31 by the information analysis module of the information server to extract different representations of information at varying levels of detail. For example, a stock quote for a company may be extracted solely by the stock price at a low detail representation, or by the stock price, the low and high of the day's price, the volume traded, and the percent change in value at a higher detail representation.

At step 32, the extracted information is classified by the information classification module of the information server according to the category and priority of the information. The information category may include weather, traffic, financial news, movies, among others. The priority of the extracted information is based on the importance of the information independent of the level of interest any individual user has in that information and on the importance of the information independent of how old the information is. The priority is calculated based on the assumption that the information has just been reported.

At step 33, the information is stored in the information database. The information database is designed to enable an Internet appliance to access the most important information for a given user's interests as well as other important inter-related information. The database records the information in categories and provides mechanisms to determine the relationships between information in different categories. These relationships define sub-categories of information and clusters of information that group related categories and sub-categories together. The clusters of information may be defined by manually programmed relationships between information categories and sub-categories, as well as by relationships formed when natural language processing of the information database can determine what information categories and sub-categories associated with the information accessed by a user would ideally be available.

Lastly, at step 34, the information rendering module extracts from the information database the information determined to be important to the user. The information extracted from the database is rendered at different levels of detail and summarization dependent on the user's interest in the information, the priority of the information, and the information rendering capabilities of the user's Internet appliance. The user's information preferences are extracted from the user's profile stored in the server. The user's profile specifies the information interests of each user in a variety of information categories, such as weather, news, stock quotes, among others. The renderings of individual items of information, which may be on a variety of categories and at different levels of detail and summarization, are combined into a concise and easy to understand summary delivered to the user at step 35 in a format and presentation optimized for the user's Internet appliance.

II. Information Retrieval

The information retrieval module of the information server reliably retrieves information from diverse Internet resources such as web servers, news feeds, and e-mail servers. To ultimately provide a robust and deterministic user interface for the user, availability of up to date Internet information must be guaranteed by the system. Standard Internet protocols do not guarantee end to end delivery of information since any number of computer or network failures can effect the delivery. Unless the user is protected from this unreliability, speed and ease of use will be compromised. The information retrieval module uses transactional queuing mechanisms with retry queues to manage requests for Internet information and handle failures. This ensures that required information content is brought into the system with the appropriate or necessary frequency to satisfy user demand.

Referring to FIG. 3, a schematic view of the transactional queuing mechanisms for retrieving Internet information is described. The information to be retrieved from the Internet by the information server is specified in a configuration object called an *information query*.

An information query object specifies the location on the Internet from which the information may be fetched, for example, the URL of a web site, the frequency during which it should be collected, for example, every hour during the day, the category of the information being collected, how to calculate the priority of the information being collected, and how the retrieved information should be parsed by the information analysis module of the information server, that is, what specific information items should be extracted from the information collected.

An information query can specify an individual item of Internet information to be collected, or more often, it can specify a collection of Internet information such as a set of stock prices from a user's portfolio or even all the stock prices on the New York Stock Exchange.

Two queues are used to manage the information requests from the Internet, normal queue 36 and retry queue 37. Normal queue 36 handles the successful information retrievals, while retry queue 37 handles the failed information retrieval attempts. Information queries are retrieved from the Internet starting at the head of normal queue 36 and finishing at the tail of normal queue 36. The same information query may be placed at different positions in normal queue 36 to retrieve the information at a particular time frequency.

A request for information query 38 at the head of normal queue 36 is sent to the Internet by the information server. There are two possible outcomes of this request. First, the information specified in information query 38 is retrieved from the Internet within a configured allowed time and so processing of the information by the information analysis and information classification modules of the information server and storage of the information in the information database can be completed.

Upon completion, information query 38 is placed at the tail of normal queue 36 as information query 39. Second, either a network failure occurs during retrieval of information query 38 or the information is not retrieved from the Internet within a configured time limit. In this case, information query 38 is placed at the tail of retry queue 37 as information query 40 to be retried at a later date.

Retry queue 37 is processed in a similar way to normal queue 36, with a request for information query 41 at the head of retry queue 37 being sent to the Internet. Information query 41 consists of a failed information request that was previously attempted by the information server in normal queue 36. Similarly to normal queue 36, there are two possible outcomes for requesting information query 41.

First, the information specified in information query 41 is retrieved from the Internet within a configured allowed time and so processing of the information by the information analysis and information classification modules of the information server and storage of the information in the information database can be completed. Upon completion, information query 41 is placed at the tail of normal queue 36 as information query 39. Second, either a network failure occurs during retrieval of information query 41 or the information is not retrieved from the Internet within a configured time limit. In this case, information query 41 is placed at the tail of retry queue 37 as information query 40 to be retried at a later date.

In such a way, normal queue 36 handles collection of information that can be handled "reliably," i.e., where information can be retrieved with a statistically acceptable frequency, and retry queue 37 handles "unreliable" Internet information. The information server uses this basic queuing mechanism to ensure transactional receipt of Internet content and efficient handling of network and Internet server failures. Advantageously, failure to retrieve information from the Internet is dealt with gracefully, that is, without data loss by the scheduling of a retry in retry queue 37.

In addition, careful configuration of time out periods on the processing of information queries allows efficient information retrievals. For example, a relatively short time out period on normal queue 36 means that slow and unreliable connections are aborted quickly, leaving normal queue 36 to continue to rapidly process other information queries. A long time out period on retry queue 37 means that the information server will wait for slower connections without holding up the receipt of the main bulk of Internet information. Further efficiency gains may be made through parallelism, both in the use of different queue pairs for different types of information content, for example, one queue pair to handle e-mails, another queue pair to handle web sites, and so on, and also by having multiple multi-threaded processors of each queue.

III. Information Analysis

Once information has been received from the Internet, it is parsed by the information analysis module of the information server to extract different representations of information at varying levels of detail. The different representations of information to be extracted are specified in the information query issued by the information retrieval module of the information server.

Referring now to FIG. 4, a schematic view of an illustrative information query is described. Information query 42 consists of an XML function that is executed by the information server to extract stock quote information from the Internet. Information query 42 contains two fields, query field 43 and information field 44. Query field 43 specifies URL 45 from which the stock price information is to be extracted, broad category 46 of the information being collected, time 47 during which the information should be collected, and how the priority of the information should be calculated (48). Upon execution of information query 42, the information server receives the source data contained in URL 45.

Information field 44 specifies how to parse the collected information, that is, what data fields (49) should be extracted from retrieved URL 45 and how. Data extraction utilizes pattern-matching techniques to extract data fields 49, which in the case of information query 42 for a stock price, include: the name of the stock, the market in which the stock is traded, the stock symbol, the last trade date and price, the change and change percent in value of the stock, the previous closing price, the volume traded on the day, among others. The process for extracting data fields 49 involves searching for known terms and markers in URL 45. These terms and markers are used as beginning and end delimiters for the data field being sought. For example, data field 50 for the stock symbol specifies data type 51, beginning delimiter 52, and end delimiter 53.

The pattern matching specification can either be built manually by a system administrator operating a user interface to highlight data fields 49 in URL 45, or automatically by using natural language analysis and heuristic techniques to recognize expected data fields in a source information of a known context. In case the information is easily accessed from a database, the natural language analysis is not required since the data fields can be directly extracted from the database. The natural language analysis is preferred to the manual method since it does not involve human interaction, which would be very slow and inefficient in retrieving Internet information during short time intervals.

IV. Information Classification

The information classification module classifies the extracted information according to its category and priority. Classifying the information according to its category begins with knowledge of where the information collected originated. In some circumstances, this may be all that is required to position the information in the correct information category. For example, if the information came from a news feed that provides only surf reports for a specific beach, then the information can be automatically categorized.

More likely, knowledge of the source information provides only broad categorization, and the final classification can only be determined in combination with data fields extracted during the information analysis. For example, a news feed providing stock quotes specifies a broad classification of any information read from the news feed as a stock quote, but this broad categorization must be combined with the company name and perhaps the market on which the stock is traded to provide the final category classification. This additional information can be determined from the data fields extracted during the information analysis, with the rules for providing the final classification category being specified in the information query. For example, information query 42 is categorized based on a combination of a broad category field and the name of the stock (54).

The information classification module calculates the priority of each item of Internet information collected based on the importance of the information independent of the level of interest any individual user has in that information and on the importance of the information independent of how old the information is. The priority is calculated based on the assumption that the information has just been reported. An individual user's interest in a category of information and the modification of priority based on how recently the information was reported are properties that are applied by the information classification module at the time at which the information is accessed by the user. The individual user's interest in a category of information is specified in the user's profile stored in the information server.

The priority of an information category is computed using arithmetic and logic equations. For example, illustrative priority equation 48 in FIG. 4 is an arithmetic equation to calculate the importance of a stock quote based on the percentage change of a stock price. Data used within the calculation is taken from the CHANGEPERCENT data field (49) extracted during information analysis.

Referring now to FIG. 5, an illustrative logic equation to compute the importance and priority of a weather report is described. Logic equation 55 determines the importance and priority of a weather report by looking for occurrences of unusual weather conditions within the weather report, such as thunderstorms, ice, sleet, wind conditions, among others. When an unusual weather pattern is present at the weather report, the priority of the weather report is increased.

The priority of an information category is represented by a numeric value such that the higher the value, the higher the priority. The priority can take values between 0 % and 100 %, and it is normalized to provide consistent values across information categories. An assignment of a priority of 95 % implies that an item of information is more important than 95 % of the information previously collected in a specific information category. In this way, historical trends can be used to normalize the calculation of priority of an item of information within a category.

It should be understood by those skilled in the art that a wide variety of equations may be used to calculate the priority of any given information category. For example, in calculating the priority of e-mail received, header information within the e-mail can be used as the basis of a calculation of the priority of an individual message. Such information can be extracted from the subject of the e-mail by looking for reply or forward markers, or from the e-mail distribution list by looking for e-mail sent directly to the user.

V. Information Database

Once the retrieved information has been analyzed and classified according to its category and priority, it is stored in an information database designed to enable an Internet appliance to access the most important information for a given user's interests as well as other important inter-related information. The information database records the information in categories and provides mechanisms to determine the relationships between information in different categories.

Referring to FIG. 6, a schematic view of an illustrative fragment of a categorization tree for managing the relationships between information items in the information database is described. Categorization tree 56 organizes the information in the information database by categories and associated relationships between information items in the categories. The categories and associated relationships are manually programmed in the information database.

Each information item in categorization tree 56 is a node in the tree, and a universally unique identifier is assigned to every node in the tree to uniquely identify a node. There are five types of nodes in categorization tree 56: (1) global root node; (2) category node; (3) cluster definition node; (4) content node; and (5) virtual content node. Each node in the tree may be of more than one type. Global root node 57 defines the head of the information space from which all of the information in the database may be traversed. Category nodes 59a and 59b represent a category of information within the information database, with category node 59a representing the entertainment category and category node 59b representing the geography category. Category node 59a is divided into category nodes 61a-c, which are sub-categories of the entertainment category, namely, movies (61a), producers (61b), and actors (61c).

Category node 61a refers to content node 63a and virtual content nodes 63b and 65. Content node 63a contains an instance of information content produced by the information analysis module of the information server, in this case, a video clip of a movie. Content nodes may be of different types, including text, graphics, imagery, video, sound, and combinations. Virtual content nodes 63b and 65 specify instances of information that vary depending on context parameters, for example, virtual content node 63b represent movie theater information that depends on location, and virtual content node 65 represent movie schedules that depend on date and time. Other virtual content nodes in categorization tree 56 include node 61d representing maps of given locations, and node 61e, representing driving directions between two locations.

In addition, category nodes 61a (movies)and virtual content node 63b (theaters) are also cluster nodes, which define a cluster of information containing related information items. Cluster node 61a define cluster 66 and cluster node 63b define cluster 67. Cluster 66 is used to specify that when information on a movie is requested by a user, related information such as the movies' producers (61b) and actors (61c), a sample preview clip of the movie (63a), and the theaters where the movie is currently playing (63b) can also be displayed to the user. Information clusters allow the users to see all relevant information on a topic and can be cascaded to form large interrelated maps of content. When a new information item belonging to an information category is inserted in the information database, any existing information cluster for that information category is checked to determine whether the new information item can be inserted in the cluster.

The nodes in categorization tree 56 are connected to each other through several arcs that define relationships between the nodes. An arc links two nodes specified by their unique identifiers, and also has properties defining the weighted strength of a relationship, the age and any expiration date of the relationship, and any ownership or privacy rules. The weighted strength of a relationship is a value associated with the arc that determines how strongly the nodes connected by the arc are related. This value is used by the information rendering module of the information server when extracting the most important information for a user from the information database. Ownership and privacy properties define which users have access to the information stored at the node destination of an arc. For example, e-mail messages are made private to a user by their ownership being associated to the user and properties being set on the arc that enforce the privacy of e-mail messages.

There are five types of arcs in a categorization tree: (1) "is a subcategory of" arc; (2) "is an instance of" arc; (3) "is equivalent to" arc; (4) "is related to"; and (5) "is a cluster definition of" arc. The "is a subcategory" arc is used to link the global root node to category nodes to form a categorization of the information space. Arcs 58a and 58b in categorization tree 56 are of this type.

The "is an instance of" arc is used to connect category nodes to content nodes or virtual content nodes to indicate that the content is in the specified category. Arcs 62a and 62b in categorization tree 56 are of this type.

The "is equivalent to" arc indicates when two or more content nodes contain equivalent content, for example, two content nodes containing news stories on the same topic and same date but from different news sources would be connected with an "is equivalent to" arc.

The "is related to" arc is used to specify relationships between instances of information in content nodes and virtual content nodes, such as arc 64 in categorization tree 56.

Lastly, the "is a cluster definition of" arc is used to connect a category node to a cluster node and that cluster node to other nodes that may contain information that if available should be presented with the requested information in the cluster node since it is closely related. Arcs 68a-d in cluster 66 of categorization tree 56 and arcs 69a-c in cluster 67 of categorization tree 56 are of this type.

Information items in the information database also can be interrelated by defining "ad hoc relationships" between them. Ad hoc relationships are formed when natural language processing of content nodes in a categorization tree can determine what other associated information are available to an user and also what information a content node can supply that would be of interest to other content nodes. This enables ad hoc relationships to be built between content nodes from anywhere within the information database without having to follow manually programmed cluster or category definitions already in place in the information database.

Referring now to FIG. 7, a schematic view of an illustrative ad hoc relationship between content nodes in a categorization tree is described. Content node 70 contains an instance of a summary plot of the movie *Antz*, and is connected to content nodes 71, 72, 73, and 74, through "is related to" arcs. Natural language processing of the information in the categorization tree has determined that content node 75 representing a TV show containing a review of the movie Antz is related to content node 70. An ad hoc relationship is then formed between content nodes 70 and 75 though "is related to" arc 77.

VI. User Configuration Interface & User Profile

A user specifies which information he or she considers to be important and relevant through a user configuration interface. The interface enables users to explicitly specify not only the categories of information of interest, but also the relative level of interest in each information category. The user configuration interface may be a user configuration web site or a call center where a user talks to a human operator. In both cases, a questionnaire about the user's interests is completed and transmitted to the information server, which translates the collected information into a user profile.

Referring to FIG. 8, a schematic view of an illustrative user profile is described. User profile 78 consists of an XML function that is executed by the information server when the user requests information from an Internet appliance. The user profile is used to extract only the information determined to be important and relevant to that user from the information database.

User profile 78 specifies user's interests in categories 79a-f, namely, e-mail (79a), financial (79b), lottery (79c), traffic (79d), weather (79e), and horoscope (79f). Category 79b has two sub-categories, 80a and 80b, further specifying the user's interest in stock prices of technology companies as the relevant financial information. Categories 79a-f have information or content items 81a-h that specify exactly what information item belonging to a given category the user has an interest in. For example, in traffic category 79d, the user only has an interest in the traffic of the San Francisco Peninsula in California, USA (81f).

Information items 81a-h have associated interest levels 82a-h that specify the importance and priority of information items 81a-h to the user. Interest levels 82a-h can be adjusted dynamically based on the user's behavior while interacting with an Internet appliance interface. For example, if information is presented to a user on a particular category of information and the user routinely skips past that information in the Internet appliance interface, the system may determine that the user has a reduced interest in this category of information. As such, the interest level on that particular category recorded in the user profile may be reduced accordingly with appropriate notification to the user.

Conversely, if a user spends a significant amount of time accessing a particular category of information in the Internet appliance interface, it can be reasonably determined that the user has a higher interest level in that information category. As such, the interest level on that particular category recorded in the user's profile can be increased accordingly with appropriate notification to the user. The dynamic adjustments to a user's interest levels are limited within reasonable bounds to avoid such adjustments having a more significant effect on the user profile than more accurate specifications of a user's interests specified in questionnaires through user configuration interfaces.

VII. Information Rendering

The user's profile is combined with the information in the information database to provide the most important information to the user through an Internet appliance interface. The information rendering module of the information server traverses an individual user's profile, taking each definition of a category of interest as a starting point within the information database, which corresponds to a category node in a categorization tree. From each of these starting points within the information database, a calculation is carried out to determine which pieces of information in the vicinity of the starting point are important enough to bring to the user's attention, that is, which content nodes and virtual content nodes can be reached from the starting point by traversing arcs in the categorization tree in the information database. The importance of a content node and virtual content node to the user and how strongly the nodes are related depend on the combined weighted strength of the arcs that are traversed.

Referring to FIG. 9, an illustrative equation to compute the importance of a content or virtual content node in a categorization tree to an user is described. Equation 83 involves three parameters, 83a-c, that are used to calculate the importance of a content or virtual content node. Equation parameter 83a consists of the multiple of all strength of relationship weights assigned to the arcs that are traversed from the starting point to get to the content or virtual node whose importance is being determined. The weights are recorded in the information database. Equation parameter 83b is extracted from the user's profile and consists of the interest level of the user in the information category represented by the category node at the starting point. Lastly, equation parameter 83c is an aging factor stored in the information database for each arc traversed from the starting point to the content or virtual node whose importance is being determined. The aging factor for a given arc is determined as the difference in dates between the date when the arc was created and the date when the user is accessing the information at the starting point.

Equation 83 is calculated for each of the user's interests so that a prioritized list of the importance of the information in the information database to be transmitted to the user is constructed. The prioritized list is used to determine the order in which the information is to be transmitted to the user when accessing an Internet appliance interface. The process of constructing the prioritized list can be optimized by using importance thresholds for arcs that are traversed from a starting point to a content or virtual content node. Arcs whose importance fall below the importance threshold will link to information items that will not be important enough to bring to the attention of the user, so that equation 83 is only calculated for the subset of arcs which is most interesting to the user.

By adjusting the importance threshold, the quantity of information pulled from the information database can be adjusted for different Internet appliance bandwidths, different Internet appliance interfaces, and different user requests. Additionally, when this process is applied to only a subset of the user's profile, a similar prioritized list can be constructed for a subset of information categories in the user's profile. These subsets allow the user to have more control of the level of detail presented for each information category when accessing an Internet appliance interface. For example, when a user initially logs into an Internet appliance interface a high level summary can be obtained of all their information interests at a relatively high importance threshold. The user then may go on to view more information, including that at a lower importance threshold in a specific information category, such as viewing the details on a stock portfolio when accessing the financial category.

The calculation of the importance of each individual item of information to a user allows a decision to be automatically made as to whether that information item should be presented to the user when accessing a particular Internet appliance. It also enables the information rendering module of the information server to decide how much detail should be provided when the information is presented (or "rendered") to the user such that an information item that is more important to the user at a moment in time is rendered with more detail than another information item that is less important. This approach helps the information rendering module of the information server present information concisely to the user.

The rendering of information is based on configured rendering definitions that are stored in the information server and specify what information is available at different levels of detail, or strata. A rendering of information at a higher strata is more summarized than a rendering of the same information at a lower strata.

Referring to FIG. 10, a schematic view of an illustrative configured rendering definition for the lottery information category is described. Rendering definition 84 consists of an XML function that is executed by the information server when information is to be rendered from the information database to an Internet appliance. Rendering definition 84 contains two rendering levels of detail, 85 and 86. Each rendering level of detail defines an information category (85a and 86a), a stratum (85b and 86b), the Internet appliance interface for which the information is being rendered (85c and 86c), and the language in which it is being rendered (85d and 86d).

Each rendering level provides formatting within which data parameters such as data parameter 87 are substituted. Data parameters substituted in renderings are extracted by the information analysis module of the information server and stored in the information database. Rendering definitions may be conditional on the values of data parameters such that detail for a particular information category is only rendered if the values of certain data parameters indicate that the user may be interested in receiving the detailed information. For example, detail on the volume of a stock may be rendered to an user if a stock volume parameter indicates that the volume in the traded stock has reached an exceptional value at a particular date.

Stratum 85b at 100 % represents a minimum rendering consisting of only a summary of the latest information on information category 85a. Stratum 86b at 50 % provides some additional detail for the user, in this case, the latest two lottery results, that may be rendered if the user's interest level in information category 86a is high enough. Any number of renderings may be configured between the 100 % stratum and 0% stratum levels. A rendering for an item of information is selected by the information rendering module of the information server when the stratum value is appropriate for the level of importance of the information to the user. The higher the importance of an item of information, the lower the stratum and the lower the summarization selected to render the information, that is, the information will be rendered with more detail.

The selection of a rendering definition depends on the Internet appliance interface (85c and 86d) for which information is being rendered. Only renderings appropriate for a user's Internet appliance are selected for rendering an item of information. The same information item may be rendered in different ways for different Internet appliances so as to optimize the effectiveness of the information on the target appliance. For example, a stock quote may be rendered in the form of a visual graph on an Internet appliance with a visual display, but on a voice-based Internet appliance the rendering may consist of only the percentage change in the stock price.

Referring to FIG. 11, a flow chart for combining renderings on a variety of information categories into a concise and easy to understand summary to be delivered to an Internet appliance is described. To make the combined rendering both easy to understand by the user and also most effective in quickly providing the most important information to an user, a two-stage sorting algorithm is carried out. First, at step 89, the list of renderings of all information categories that are important to the user are sorted so that the most important information categories are rendered to the user before the less important ones. Second, at step 90, a sort is performed within each information category to order the individual information items in each information category by their importance to a user. For example, ordering individual information items in the stock quotes category would render a stock price that has risen substantially prior to one that is stable.

Providing a concise and easy to understand summary optimized for network bandwidth and display capabilities of each individual Internet appliance optionally involves a method for transforming text so that it can be presented using a voice-based interface.

Referring now to FIG. 12, a flow chart for transforming written text so that it may be presented by a voice interface as it is normally spoken in conversation is described. At step 93, an information context analysis is performed so that the broad context of the information is determined by the information rendering module. The broad context may determine, for example, that the information item is an e-mail message, or that the information item contains telephone numbers of a movie theater.

The context analysis searches for known keywords within the information rendering that indicate the context of the information, for example, a paragraph of information that contains the words "highway" and "interstate" is likely to be discussing roads, a sentence that starts with the word "phone" is likely to contain a phone number, a sentence that contains the word "dollar" or the character "$" is likely to contain financial figures, and so on. At step 94, a word token analysis is performed by the information rendering module to determine word tokens whose understandability when spoken in conversation can typically be improved. Examples of such word tokens include streams of numbers or dates, or header information in an e-mail message.

By combining steps 93 and 94, a high degree of reliability can be achieved in improving the understandability of text. Referring now to FIG. 13, a schematic view of illustrative written text to voice transformations for voice-based Internet appliance interfaces is described. Information context analysis 96, token analysis 97, and improved renderings 98 contain examples of transformations of different information items in particular contexts, including information items in financial context 99 and information items in telephone context 100.

Information items in financial context 99 are rendered in such a way to clarify specific symbols used in a financial context, such as "$" signs and "½" fractions of stock quotes, which are rendered in speech as the word "dollar" and the word "half", respectively. Similarly, telephone numbers in telephone context 100 are rendered so that well-known phone prefixes, such as "1-800" and area codes, are easily understood in a voice-based interface.

VIII. Internet Appliance Interface

The rendered information can be accessed from an Internet appliance through an Internet appliance interface, that may be a voice interface or a visual interface, depending on the display capabilities of the Internet appliance being used. The Internet appliance interface is designed to that information can be accessed from any Internet appliance with little user interface interaction. The Internet appliance interface preferably provides the a high degree of desired Internet information to an user in a conveniently accessible manner from any Internet appliance.

Referring to FIG. 14, an example text transcript of user access to an illustrative voice-based Internet appliance interface is described. Text transcript 101 contains the text translation of voice prompts issued by the user of a voice-based Internet appliance interface as well as the text translation of the voice information sent by the information server to the user during a user's phone call to the server. Upon dialing the phone number corresponding to the server, the user receives summary 102 of the information important to the user. After a brief pause following summary 102, the information server presents the user menu of options 103 to prompt the user for access to the user's pre-selected information categories.

When the user selects an information category option by saying the number corresponding to the selected information category, the server fetches the appropriate rendering configuration of the information, and translates the information to the user through information voice notification 104. The information transmitted to the user for each information item depends on the importance of the information item, with more detail information being provided for the information items that are considered more important.

A brief pause follows information voice notification 104, and the user then hears menu 103 again to select further information categories. The user selects the e-mail information category, and listens to e-mail notification 105 from the server. Selecting the e-mail information category again after listening to e-mail notification 105 triggers the server to render the important messages in e-mail notification 106 providing more detail to the user. At any point during the phone call the user may choose to hangup, or to return to menu of options 103 to access more detailed information on the relevant information categories.

Referring to FIGS. 15A and 15B, schematic views of illustrative visual-based interfaces for use with an Internet appliance having a visual display is described. Visual-based interface 107 shown in FIG. 15A contains top-level summaries on various information categories of interest to an user named Dave, including e-mail messages, financial news, weather information, shopping requests, news headlines, and entertainment news. Each information category is easily accessed through icons displayed in visual-based interface 107. Accessing the information category icons enables the user to get more detail information on any one of the relevant information categories. Visual-based interface 108 shown in FIG. 15B consists of an instance of user interaction requesting information on the movie *Antz*. Visual-based interface 108 displays all the related information on the movie *Antz* that is available on the information database.

Although particular embodiments of the present invention have been described above in detail, it will be understood that this description is merely for purposes of illustration. Specific features of the invention are shown in some drawings and not in others, and this is for convenience only and any feature may be combined with another in accordance with the invention. Steps of the described processes may be reordered or combined, and other steps may be included. Further variations will be apparent to one skilled in the art in light of this disclosure and are intended to fall within the scope of the appended claims.

What Is Claimed Is:

1. A system for providing a user with access to items of Internet information according to a user's preferences using an Internet appliance, the system comprising:

an information server programmed to communicate with an Internet appliance;

an information database programmed to store and manage Internet information;

an user configuration interface to enable an user to generate a user profile to specify the user's preferences; and

an Internet appliance interface to enable an user to access Internet information using the Internet appliance.

2. The system of claim 1, wherein the Internet appliance comprises an electronic device configured to access the Internet.

3. The system of claim 1, wherein the information server comprises:

an information retrieval software module;

an information analysis software module;

an information classification software module; and

an information rendering software module.

4. The system of claim 2, wherein the information retrieval software module comprises information retrieval software routines for retrieving Internet information from a plurality of Internet resources.

5. The system of claim 4, wherein the Internet resources comprise web servers, news feeds, and e-mail servers.

6. The system of claim 4, wherein the information retrieval software routines further comprise transactional queues that manage requests for Internet information.

7. The system of claim 3, wherein the information analysis module comprises:

software routines for parsing retrieved Internet information and extracting information representations at various levels of detail; and

pattern matching software routines for searching and extracting known markers and data fields from retrieved Internet information.

8. The system of claim 7, wherein software routines for parsing retrieved Internet information and extracting information representations comprise information queries for specifying one or more of: information representations to be extracted from the retrieved Internet information; a location from which Internet information is to be extracted; a category of the information being extracted; a time during which the information is extracted; and how to determine the priority of the information being extracted.

9. The system of claim 3, wherein the information classification module comprises software routines for classifying retrieved Internet information according to information categories and information priority.

10. The system of claim 9, wherein the information categories comprise one or more of: a messages categories; a news category; a financial category; a weather category; an entertainment category; a traffic category; and a calendar category.

11. The system of claim 9, wherein the information priority is computed based on the importance of the information independent of a level of user interest and independent of an age of the information.

12. The system of claim 3, wherein the information rendering module comprises:

software routines for extracting the information from the information database determined to be desired by the user;

software routines for rendering items of information at one or more levels of detail; and

software routines for combining renderings of items of information into a summary format.

13. The system of claim 12, wherein software routines for extracting the information comprise software routines for computing the importance of an item of information in the information database based on related information items in the information database and on a level of user interest.

14. The system of claim 13, wherein the level of user interest comprises a priority value associated with the item of information item or information category stored in the user's profile.

15. The system of claim 13, wherein software routines for rendering individual information items at one or more levels of detail comprise software routines for extracting the items of information from the information database specified in a plurality of rendering definitions stored in the information server.

16. The system of claim 15, wherein the rendering definitions comprise formatting specifications for representing the items of information at a plurality of levels of detail.

17. The system of claim 12, wherein software routines for combining renderings of items of information items into a summary format comprise:

a routine for sorting the renderings of the items of information and information categories that are determined to be important to the user; and

a routine for transforming the items of information into a format for presentation by a voice interface.

18. The system of claim 17, wherein the items of information and information categories that are determined to be important to the user are specified in the user's profile with a high priority value.

19. The system of claim 17, wherein routines for transforming the items of information comprises information context analysis software routines and word token analysis software routines.

20. The system of claim 1, wherein the information database comprises software routines for recording items of information in categories and determining relationships between items of information in different categories.

21. The system of claim 20, wherein the software routines organize the items of information in a categorization tree.

22. The system of claim 21, wherein the categorization tree comprises nodes that identify an information item or category and arcs that define relationships between the nodes.

23. The system of claim 22, wherein the nodes comprise a universally unique identifier.

24. The system of claim 1, wherein the user configuration interface comprises a user configuration web site or a human-operated call center.

25. An information server programmed to communicate with an Internet appliance, the information server comprising:

an information retrieval software module;

an information analysis software module;

an information classification software module; and

an information rendering software module.

26. The server of claim 25, wherein the information retrieval software module comprises information retrieval software routines for retrieving

Internet information from a plurality of Internet resources.

27. The server of claim 25, wherein the information analysis software module comprises:

software routines for parsing retrieved Internet information and extracting information representations at a pre-determined level of detail; and

pattern matching software routines for searching and extracting known markers and data fields from retrieved Internet information.

28. The server of claim 25, wherein the information classification module comprises software routines for classifying retrieved Internet information according to information categories and information priority.

29. The server of claim 25, wherein the information rendering module comprises:

software routines for extracting information from an information database determined to be important to a user;

software routines for rendering items of information at one or more levels of detail; and

software routines for combining renderings of the items of information into a summary format.

30. The server of claim 29, wherein the information database comprises a database of Internet information and software routines for recording the information in categories and determining relationships between items of information in different categories.

31. A method for providing Internet information using an Internet appliance according to a user's preferences, the method comprising:

providing a user configuration interface to enable a user to specify his or her information preferences;

retrieving information from the Internet and organizing the information into an information database;

extracting information from the information database determined to be desired by the user;

summarizing the extracted information at a level of detail according to a network bandwidth and display capability of the Internet appliance; and

providing an Internet appliance interface for accessing the extracted information.

32. The method of claim 31, wherein providing a user configuration interface comprises providing a web site or a human-operated call center for the user to specify information preferences.

33. The method of claim 31, wherein retrieving information from the Internet and organizing the information into an information database comprises:

retrieving Internet information from a plurality of Internet resources;

parsing the retrieved Internet information and extracting information representations at various levels of detail;

classifying the information representations into information categories;

assigning a priority level to the information categories; and

recording the information in the information in the information database according to categories and relationships between information items in different categories.

34. The method of claim 31, wherein extracting information from the information database comprises computing a level of importance of an item of information based on related items of information database and on a user's level of interest in the item of information or information category comprising the item of information.

35. The method of claim 31, wherein summarizing the extracted information comprises providing formatting specifications for representing items of information at a plurality of levels of detail and sorting the items of information.

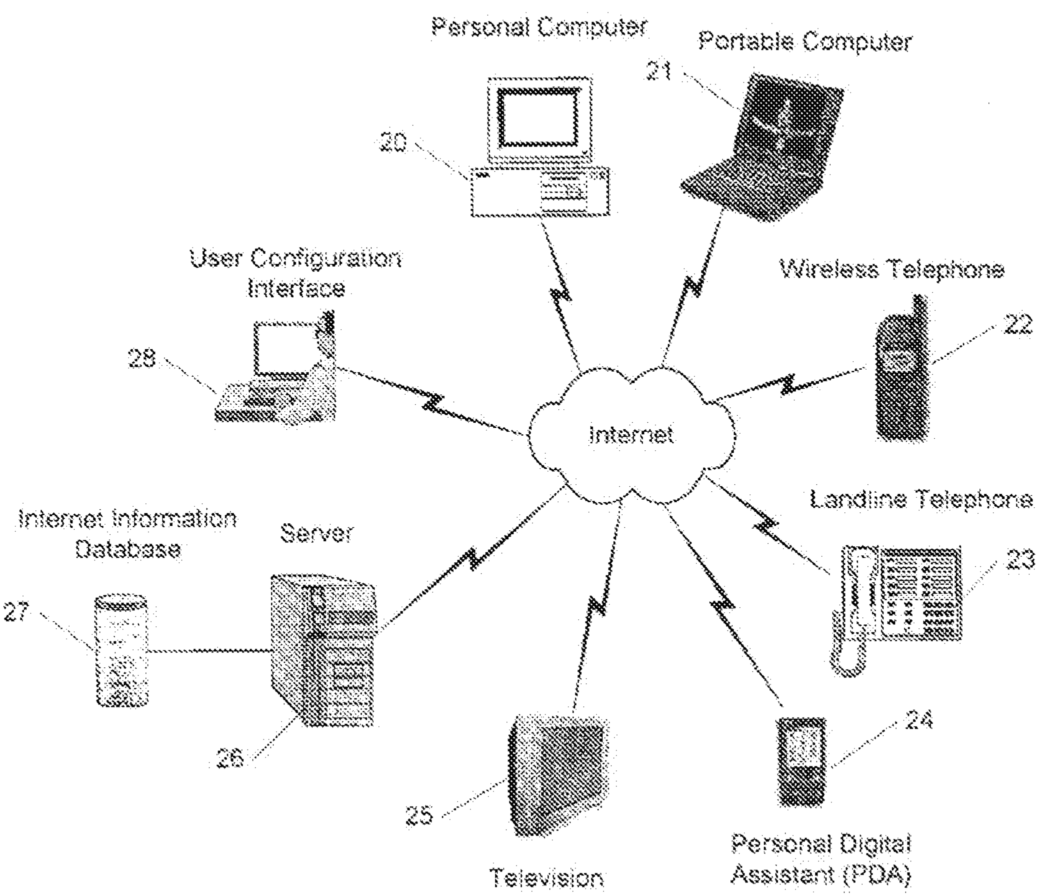
Personal Computer
Portable Computer
21
20
User Configuration Interface
Wireless Telephone
22
28
Internet
Landline Telephone
Internet Information Database
Server
23
27
24
26
25
Personal Digital Assistant (PDA)
Television

FIG. 1

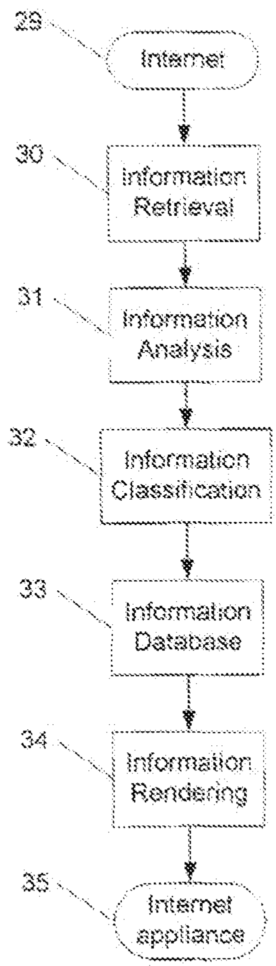
29
Internet
30
Information Retrieval
31
Information Analysis
32
Information Classification
33
Information Database
34
Information Rendering
35
Internet appliance

FIG. 2

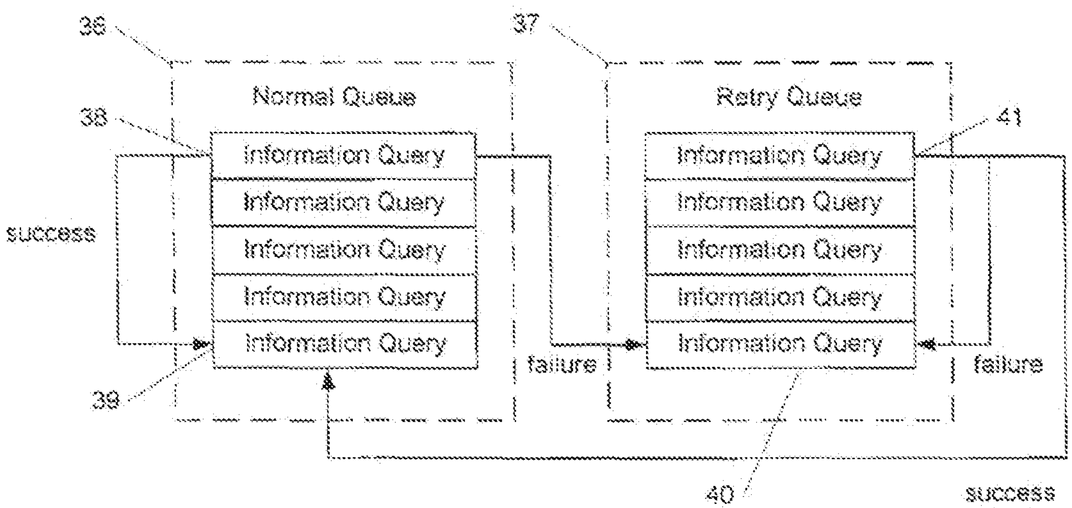
36
37
Normal Queue
Retry Queue
38
41
Information Query
Information Query
Information Query
Information Query
Information Query
Information Query
Information Query
Information Query
Information Query
Information Query
success
failure
failure
39
40
success

FIG. 3

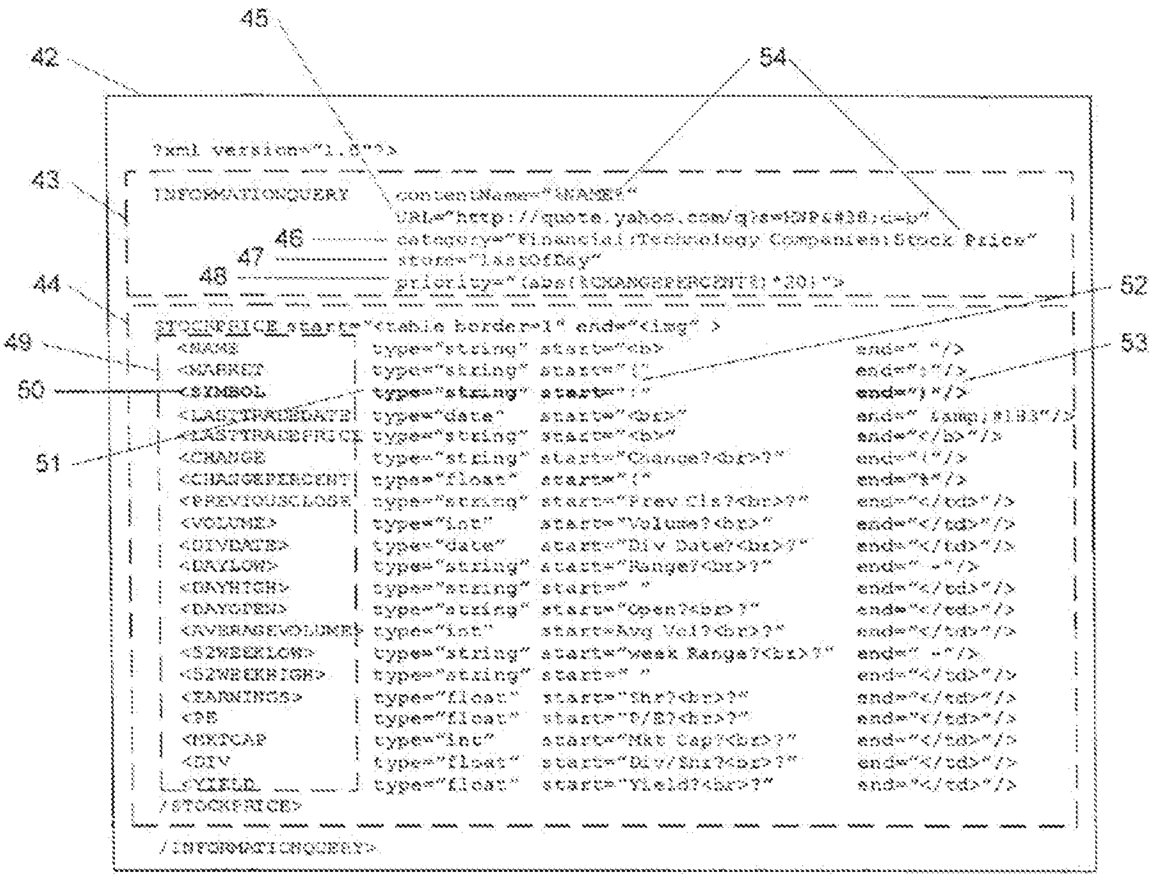
45
42
54
43
46
47
48
44
52
49
53
50
51
?xml version="1.0"?>
INFORMATIONQUERY contentName="%NAME%"
URL="http://quote.yahoo.com/q?s=%SP&d=b"
category="Financial:Technology Companies:Stock Price"
store="LastOfDay"
priority="(abs(%CHANGEPERCENT%)*20);">
STOCKPRICE start="<table border=1" end="<img" >
<NAME type="string" start="" end=" "/>
<MARKET type="string" start="(" end=":"/>
<SYMBOL type="string" start=":" end=")"/>
<LASTTRADEDATE type="date" start="
" end=" &#183"/>
<LASTTRADEPRICE type="string" start="" end=""/>
<CHANGE type="string" start="Change?
?" end="("/>
<CHANGEPERCENT type="float" start="(" end="%"/>
<PREVIOUSCLOSE type="string" start="Prev Cls?
?" end="</td>"/>
<VOLUME> type="int" start="Volume?
" end="</td>"/>
<DIVDATE> type="date" start="Div Date?
?" end="</td>"/>
<DAYLOW> type="string" start="Range?
?" end=" -"/>
<DAYHIGH> type="string" start=" " end="</td>"/>
<DAYOPEN> type="string" start="Open?
?" end="</td>"/>
<AVERAGEVOLUME> type="int" start=Avg Vol?
?" end="</td>"/>
<52WEEKLOW> type="string" start="week Range?
?" end=" -"/>
<52WEEKHIGH> type="string" start=" " end="</td>"/>
<EARNINGS> type="float" start="Shr?
?" end="</td>"/>
<PE type="float" start="P/E?
?" end="</td>"/>
<MKTCAP type="int" start="Mkt Cap?
?" end="</td>"/>
<DIV type="float" start="Div/Shr?
?" end="</td>"/>
<YIELD type="float" start="Yield?
?" end="</td>"/>
/STOCKPRICE>
/INFORMATIONQUERY>

FIG. 4

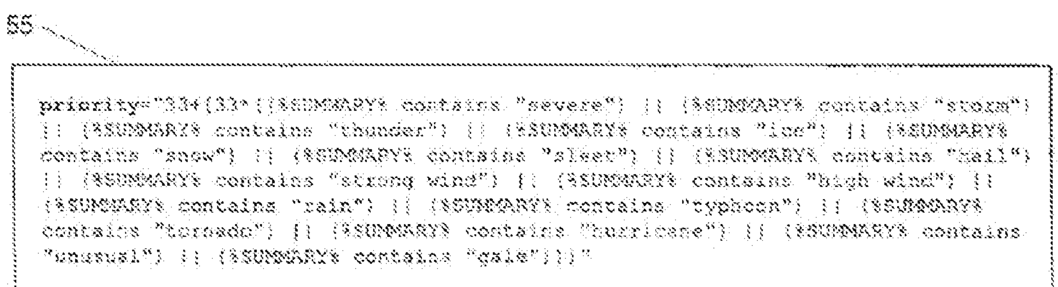
55
priority="33+(33*((%SUMMARY% contains "severe") || (%SUMMARY% contains "storm")
|| (%SUMMARY% contains "thunder") || (%SUMMARY% contains "ice") || (%SUMMARY%
contains "snow") || (%SUMMARY% contains "sleet") || (%SUMMARY% contains "hail")
|| (%SUMMARY% contains "strong wind") || (%SUMMARY% contains "high wind") ||
(%SUMMARY% contains "rain") || (%SUMMARY% contains "typhoon") || (%SUMMARY%
contains "tornado") || (%SUMMARY% contains "hurricane") || (%SUMMARY% contains
"unusual") || (%SUMMARY% contains "gale")))"

FIG. 5

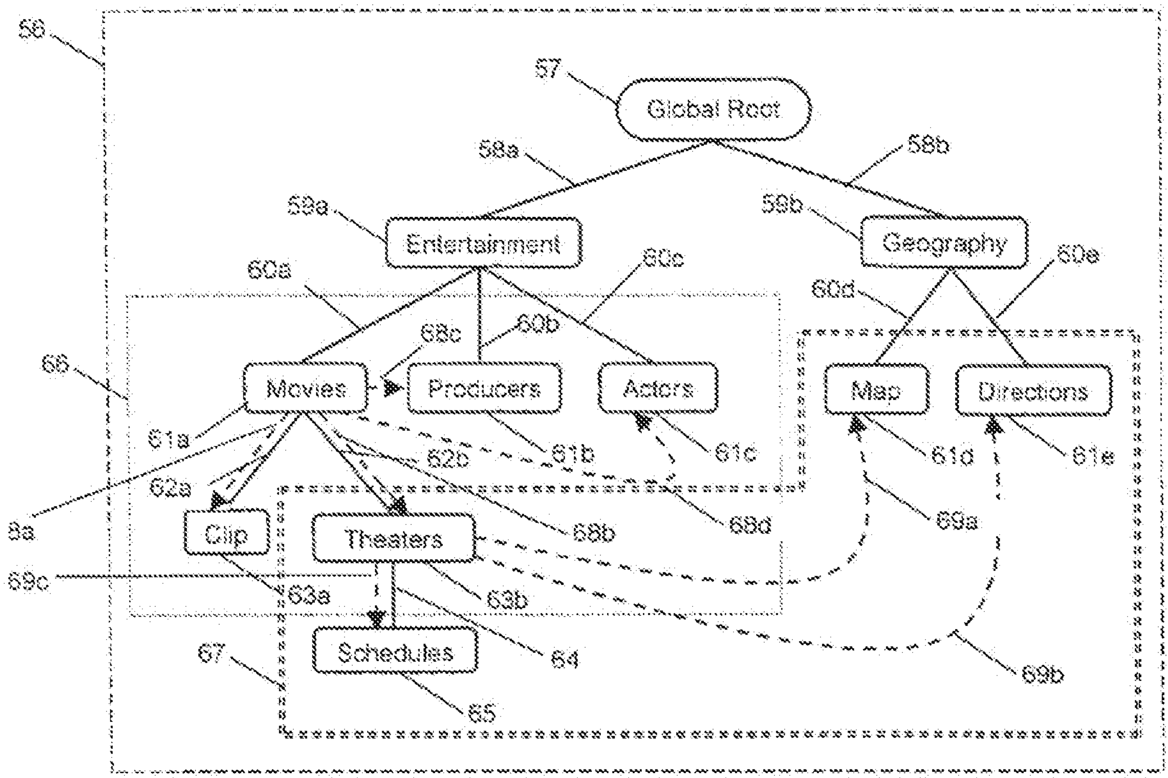
56
57
Global Root
58a
58b
59a
Entertainment
59b
Geography
60a
60c
60e
60d
60b
68c
66
Movies
Producers
Actors
Map
Directions
61a
62b
61b
61c
61d
61e
62a
8a
Clip
Theaters
68b
68d
69a
69c
63a
63b
67
Schedules
64
69b
65

FIG. 6

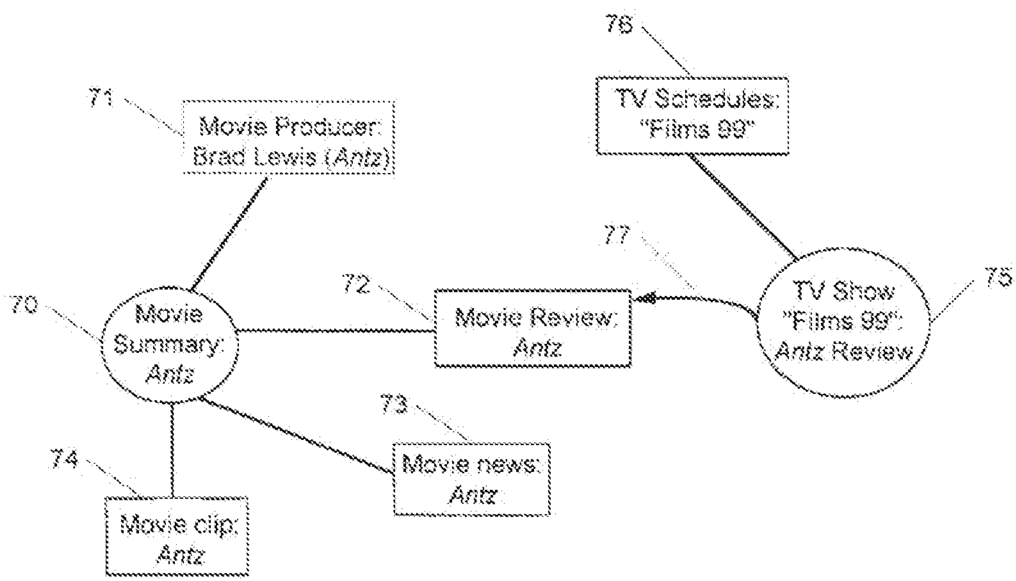
76
TV Schedules: "Films 99"
71
Movie Producer: Brad Lewis (*Antz*)
77
72
Movie Review: *Antz*
TV Show "Films 99": *Antz* Review
75
70
Movie Summary: *Antz*
73
Movie news: *Antz*
74
Movie clip: *Antz*

FIG. 7

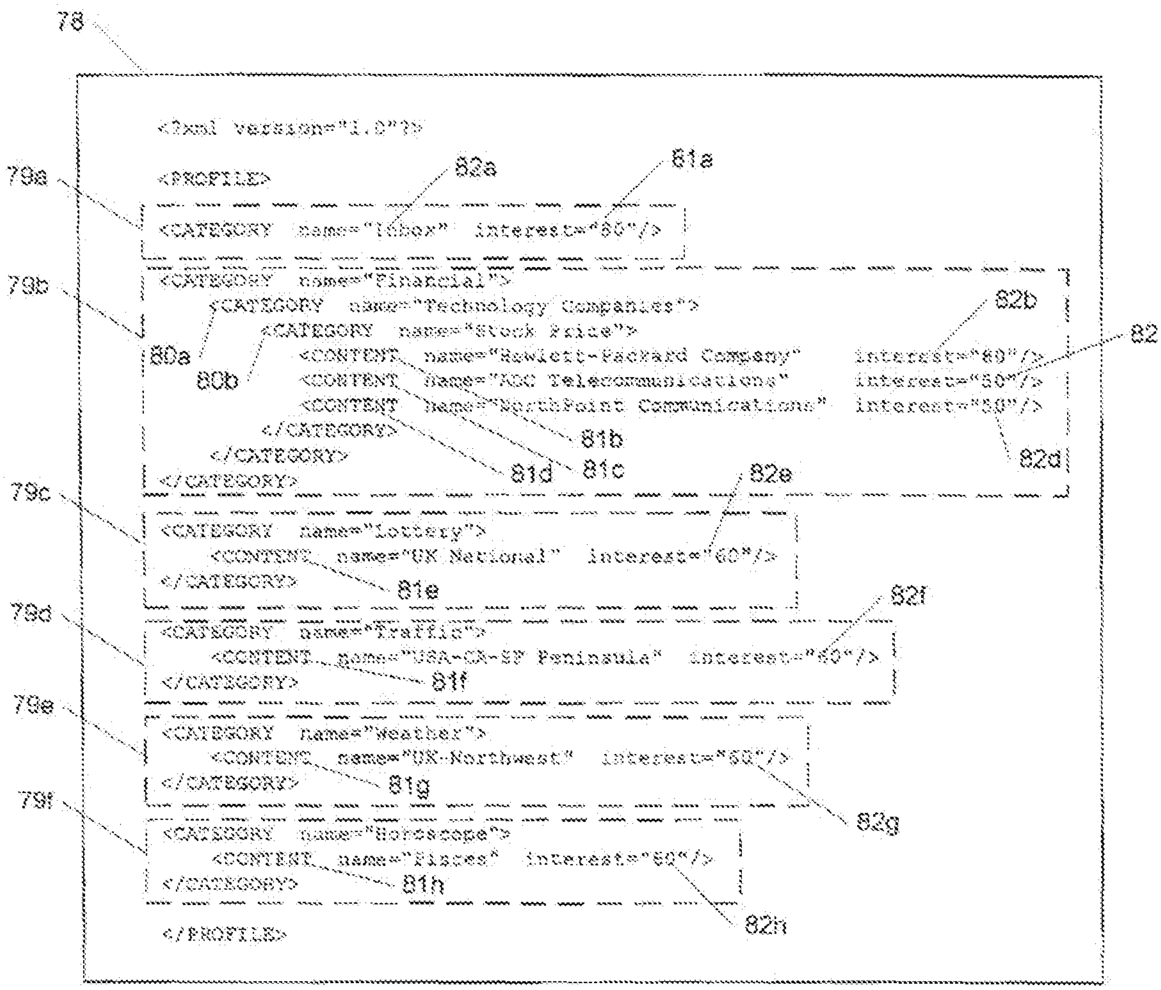
78
<?xml version="1.0"?>
<PROFILE>
79a
82a
81a
<CATEGORY name="Inbox" interest="80"/>
79b
<CATEGORY name="Financial">
<CATEGORY name="Technology Companies">
<CATEGORY name="Stock Price">
<CONTENT name="Hewlett-Packard Company" interest="80"/>
<CONTENT name="XOC Telecommunications" interest="50"/>
<CONTENT name="NorthPoint Communications" interest="50"/>
</CATEGORY>
</CATEGORY>
</CATEGORY>
80a
80b
82b
82
81b
81c
81d
82d
79c
82e
<CATEGORY name="Lottery">
<CONTENT name="UK National" interest="60"/>
</CATEGORY>
81e
79d
82f
<CATEGORY name="Traffic">
<CONTENT name="USA-CA-SF Peninsula" interest="60"/>
</CATEGORY>
81f
79e
<CATEGORY name="Weather">
<CONTENT name="UK-Northwest" interest="60"/>
</CATEGORY>
81g
82g
79f
<CATEGORY name="Horoscope">
<CONTENT name="Pisces" interest="60"/>
</CATEGORY>
81h
82h
</PROFILE>

FIG. 8

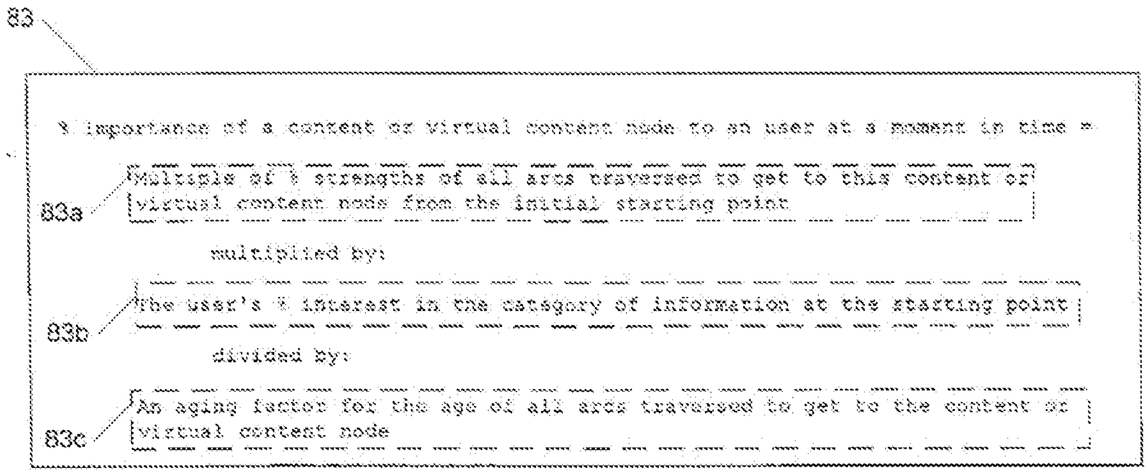
83
% importance of a content or virtual content node to an user at a moment in time =
83a
Multiple of % strengths of all arcs traversed to get to this content or virtual content node from the initial starting point
multiplied by:
83b
The user's % interest in the category of information at the starting point
divided by:
83c
An aging factor for the age of all arcs traversed to get to the content or virtual content node

FIG. 9

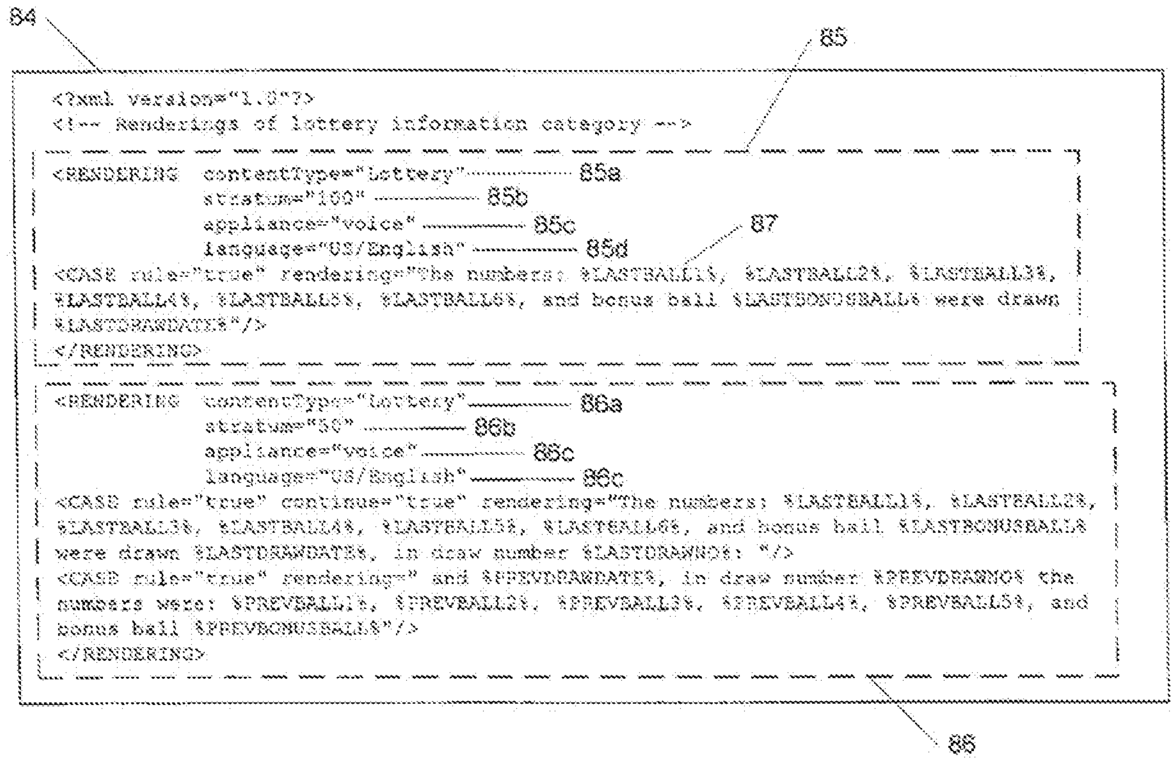
84
85
<?xml version="1.0"?>
<!-- Renderings of lottery information category -->
<RENDERING contentType="Lottery" 85a
stratum="100" 85b
appliance="voice" 85c
language="US/English" 85d
87
<CASE rule="true" rendering="The numbers: %LASTBALL1%, %LASTBALL2%, %LASTBALL3%, %LASTBALL4%, %LASTBALL5%, %LASTBALL6%, and bonus ball %LASTBONUSBALL% were drawn %LASTDRAWDATE%"/>
</RENDERING>
<RENDERING contentType="Lottery" 86a
stratum="50" 86b
appliance="voice" 86c
language="US/English" 86c
<CASE rule="true" continue="true" rendering="The numbers: %LASTBALL1%, %LASTBALL2%, %LASTBALL3%, %LASTBALL4%, %LASTBALL5%, %LASTBALL6%, and bonus ball %LASTBONUSBALL% were drawn %LASTDRAWDATE%, in draw number %LASTDRAWNO%: "/>
<CASE rule="true" rendering=" and %PREVDRAWDATE%, in draw number %PREVDRAWNO% the numbers were: %PREVBALL1%, %PREVBALL2%, %PREVBALL3%, %PREVBALL4%, %PREVBALL5%, and bonus ball %PREVBONUSBALL%"/>
</RENDERING>
86

FIG. 10

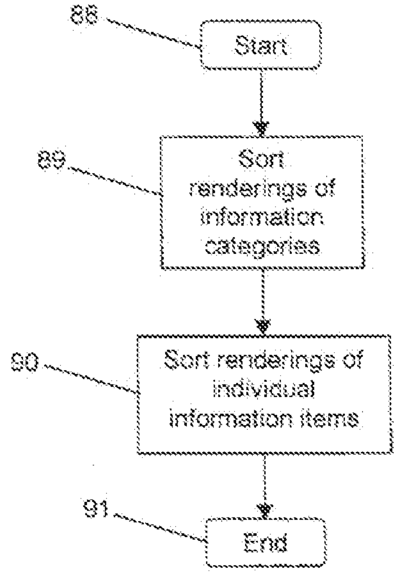
88
Start
89
Sort renderings of information categories
90
Sort renderings of individual information items
91
End

FIG. 11

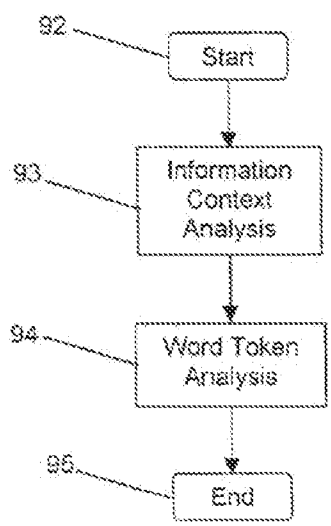
92
Start
93
Information Context Analysis
94
Word Token Analysis
95
End

FIG. 12

96 97 98 99 100

Context Analysis	Token Analysis	Improved Rendering
Financial Context	$1M	1 million dollars
Financial Context	$12 ½	12 and a half dollars
Date Context	2-3 February	2^{nd} to 3^{rd} of February
Date Context	10/13/2000	'yesterday' or 'next Saturday' or 'last Wednesday' as appropriate
Telephone Context	1-800-832-5843	1 eight hundred 8-3-2 5-8-4-3
Telephone Context	408-543-6500	4-o-8 5-4-3 6-5 hundred
Highway Context	I880	Interstate 8-eighty
Highway Context	101	1-o-1
e-mail Context	Attached message header	'An attached message sent by Joe Smith yesterday'
e-mail Context	☺	Smiley face
e-mail Context	nicks@sirenic.com	e-mail address: nicks at sirenic dot com
Any Context	http://www.sirenic.com	web address: www dot sirenic dot com

FIG. 13

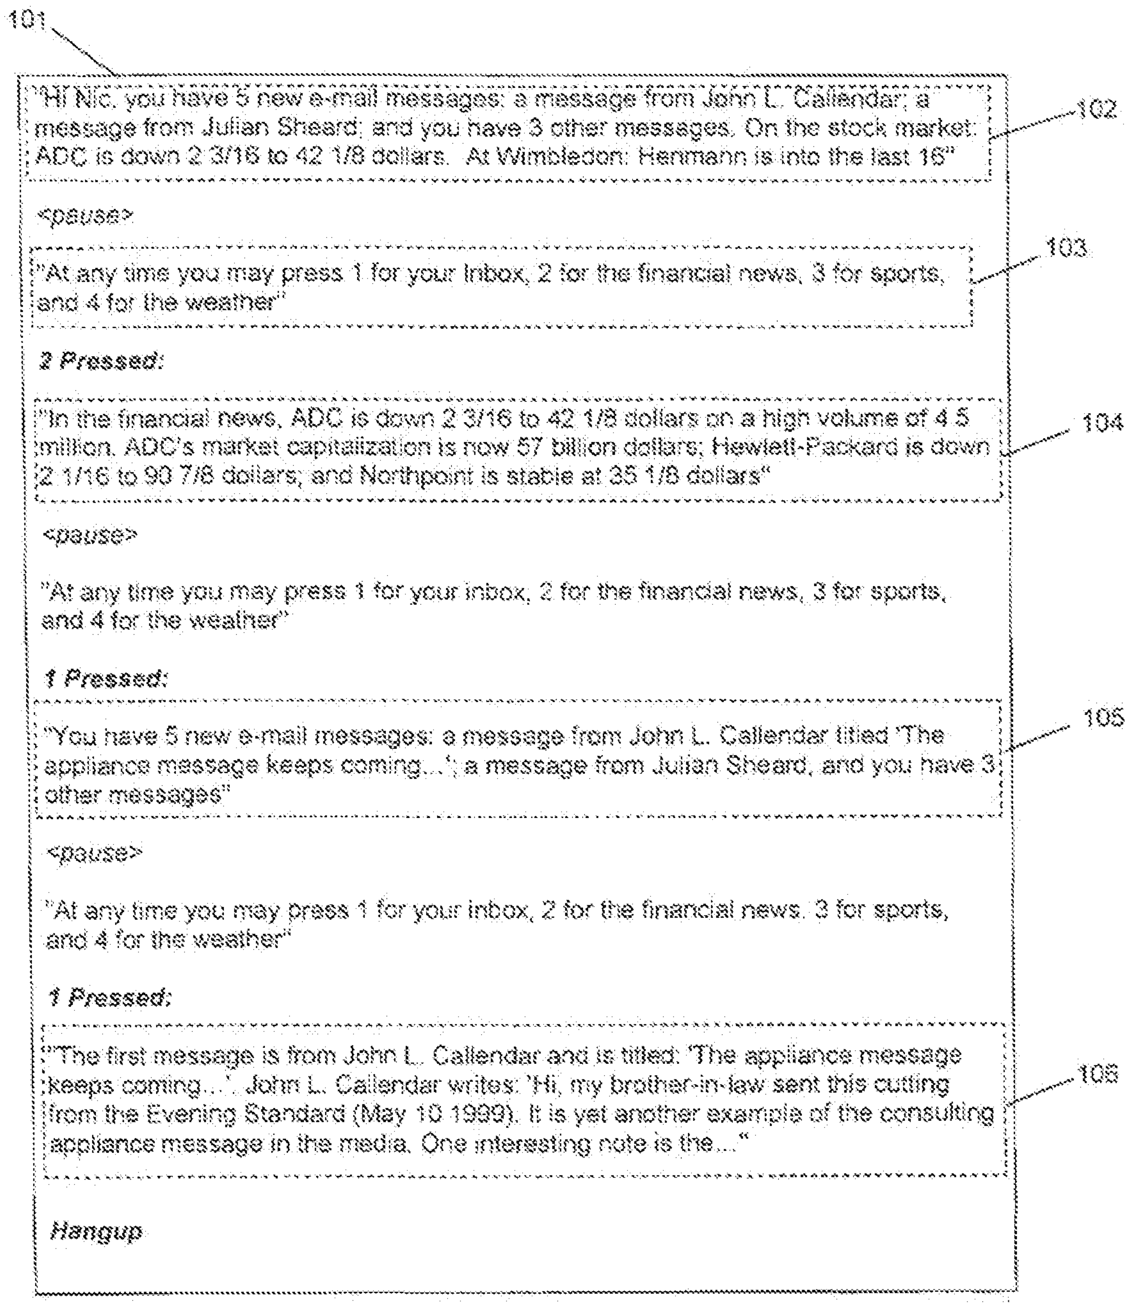
101
"Hi Nic, you have 5 new e-mail messages: a message from John L. Callendar; a message from Julian Sheard; and you have 3 other messages. On the stock market: ADC is down 2 3/16 to 42 1/8 dollars. At Wimbledon: Henmann is into the last 16"
102
<pause>
"At any time you may press 1 for your Inbox, 2 for the financial news, 3 for sports, and 4 for the weather"
103
2 Pressed:
"In the financial news, ADC is down 2 3/16 to 42 1/8 dollars on a high volume of 4.5 million. ADC's market capitalization is now 57 billion dollars; Hewlett-Packard is down 2 1/16 to 90 7/8 dollars; and Northpoint is stable at 35 1/8 dollars"
104
<pause>
"At any time you may press 1 for your inbox, 2 for the financial news, 3 for sports, and 4 for the weather"
1 Pressed:
"You have 5 new e-mail messages: a message from John L. Callendar titled 'The appliance message keeps coming...'; a message from Julian Sheard, and you have 3 other messages"
105
<pause>
"At any time you may press 1 for your inbox, 2 for the financial news, 3 for sports, and 4 for the weather"
1 Pressed:
"The first message is from John L. Callendar and is titled: 'The appliance message keeps coming...'. John L. Callendar writes: 'Hi, my brother-in-law sent this cutting from the Evening Standard (May 10 1999). It is yet another example of the consulting appliance message in the media. One interesting note is the..."
106
Hangup

FIG. 14

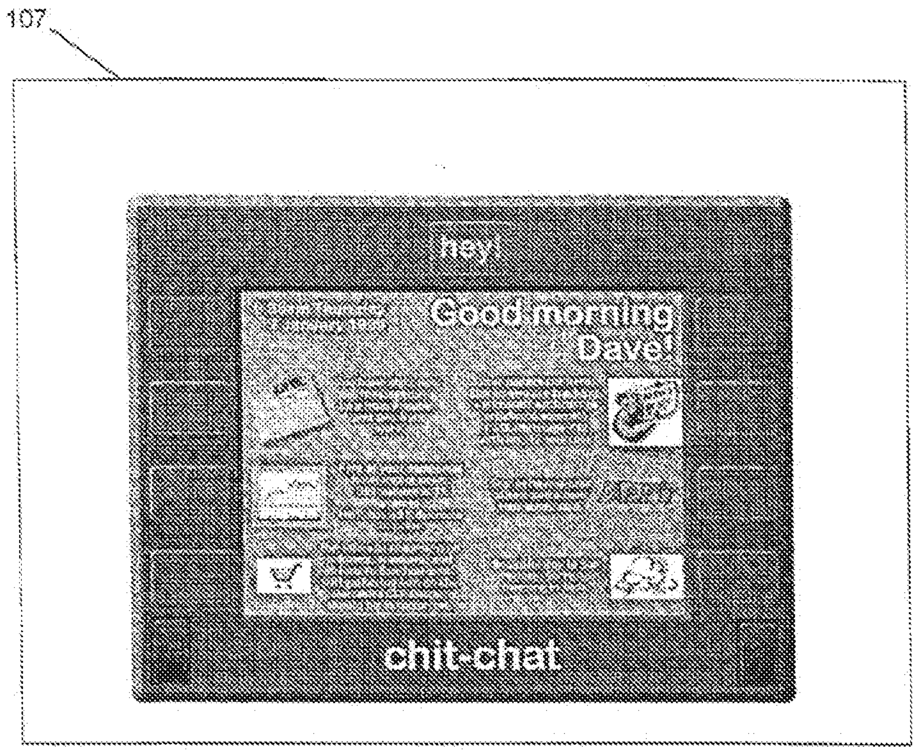

FIG. 15A

108

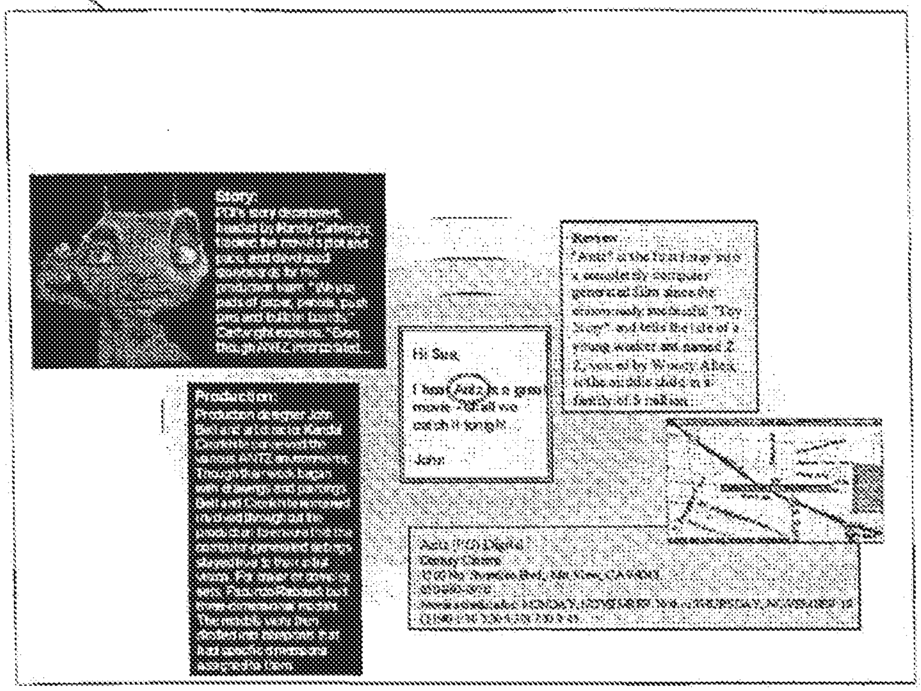

FIG. 15B

(19) Canadian Intellectual Property Office
An Agency of Industry Canada

Office de la Propriété Intellectuelle du Canada
Un organisme d'Industrie Canada

(11) CA 2 553 120 (13) A1

(40) 23.06.1994
(43) 23.06.1994

(12)

(21) 2 553 120

(51) Int. Cl.:

(22) 02.12.1993

(30) 07/991,074 US 09.12.1992

(60) 22.08.2006
(62) 2447895

(71)
SEDNA PATENT SERVICES, LLC.
1500 Market Street
27th Floor, West Tower, PHILADELPHIA, XX (US).

(72)
HENDRICKS, JOHN S. (US).
BONNER, ALFRED E. (US).

(74)
RICHES, MCKENZIE & HERBERT LLP

(54) CONTROLEUR DE RESEAU POUR SYSTEMES DE TELEVISION CABLES
(54) NETWORK CONTROLLER FOR CABLE TELEVISION DELIVERY SYSTEMS

(57)
A novel network controller for use with a digital cable headend capable of monitoring and controlling set top terminals in a television program delivery system is described. The invention relates to methods and apparatus for a network controller that manages a configuration of set top terminals in a program delivery system. The invention is particularly useful in program delivery systems with hundreds of channels of programming, a menu driven program selection system, and a program control information signal that carries data and identifies available program choices. Specifically, the invention modifies a program control information signal at the cable headend before the modified signal is transmitted to each set top terminal. This signal is used with polling methods to receive upstream data from the set top terminals. The invention initiates such upstream data retrieval, gathers all data received and compiles viewer demographics information and programs watched information. The invention processes this data and information to generate packages of advertisements, as well as account and billing reports, targeted towards each set top terminal. The invention uses upstream data reception hardware, databases and processing hardware and software to accomplish these functions.

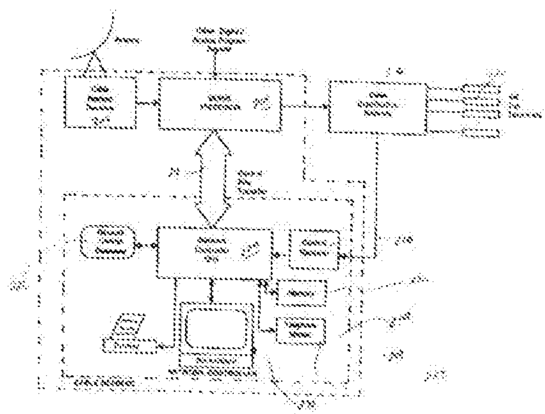

Office de la Propriété Intellectuelle du Canada

Un organisme d'Industrie Canada

Canadian Intellectual Property Office

An agency of Industry Canada

CA 2553120 A1 1994/06/23

(21) 2 553 120

(12) DEMANDE DE BREVET CANADIEN
CANADIAN PATENT APPLICATION

(13) A1

(22) Date de dépôt/Filing Date: 1993/12/02

(41) Mise à la disp. pub./Open to Public Insp.: 1994/06/23

(62) Demande originale/Original Application: 2 447 895

(30) Priorité/Priority: 1992/12/09 (US07/991,074)

(51) Cl.Int./Int.Cl. *H04N 7/16* (2006.01), *H04H 9/00* (2006.01)

(71) Demandeur/Applicant: SEDNA PATENT SERVICES, LLC, US

(72) Inventeurs/Inventors:
HENDRICKS, JOHN S., US;
BONNER, ALFRED E., US

(74) Agent: RICHES, MCKENZIE & HERBERT LLP

(54) Titre : CONTROLEUR DE RESEAU POUR SYSTEMES DE TELEVISION CABLES
(54) Title: NETWORK CONTROLLER FOR CABLE TELEVISION DELIVERY SYSTEMS

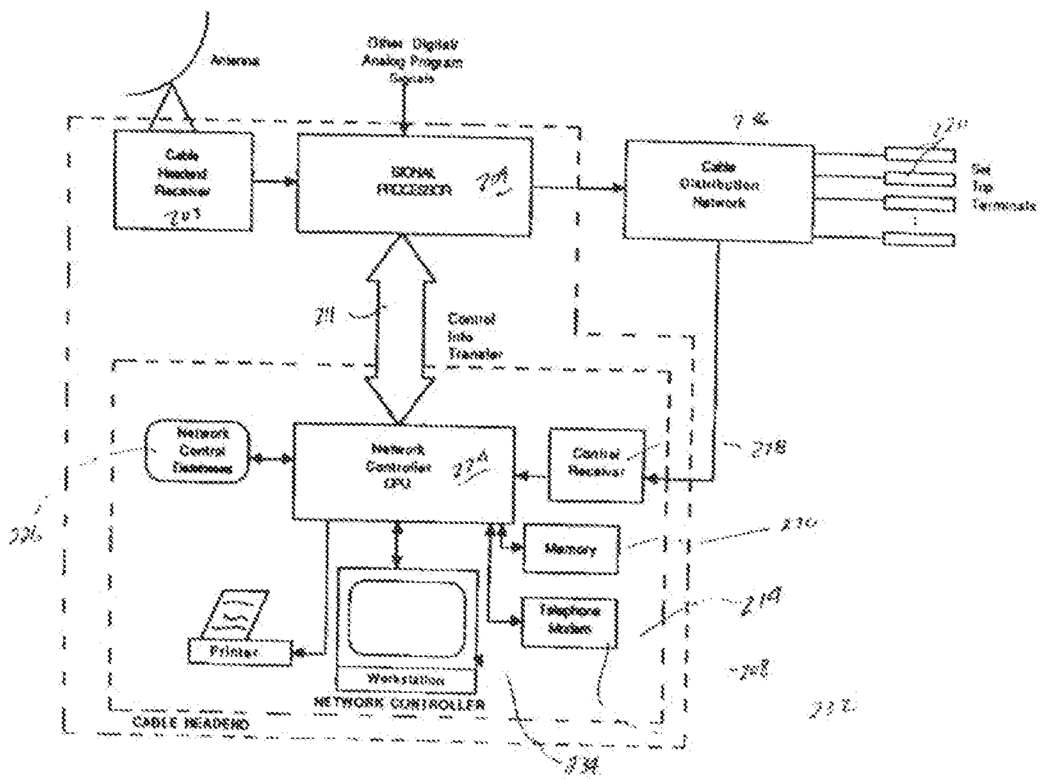

(57) Abrégé/Abstract:
A novel network controller for use with a digital cable headend capable of monitoring and controlling set top terminals in a television program delivery system is described. The invention relates to methods and apparatus for a network controller that manages a

(57) Abrégé(suite)/Abstract(continued):
configuration of set top terminals in a program delivery system. The invention is particularly useful in program delivery systems with hundreds of channels of programming, a menu driven program selection system, and a program control information signal that carries data and identifies available program choices. Specifically, the invention modifies a program control information signal at the cable headend before the modified signal is transmitted to each set top terminal. This signal is used with polling methods to receive upstream data from the set top terminals. The invention initiates such upstream data retrieval, gathers all data received and compiles viewer demographics information and programs watched information. The invention processes this data and information to generate packages of advertisements, as well as account and billing reports, targeted towards each set top terminal. The invention uses upstream data reception hardware, databases and processing hardware and software to accomplish these functions.

ABSTRACT

A novel network controller for use with a digital cable headend capable of monitoring and controlling set top terminals in a television program delivery system is described. The invention relates to methods and apparatus for a network controller that manages a configuration of set top terminals in a program delivery system. The invention is particularly useful in program delivery systems with hundreds of channels of programming, a menu driven program selection system, and a program control information signal that carries data and identifies available program choices. Specifically, the invention modifies a program control information signal at the cable headend before the modified signal is transmitted to each set top terminal. This signal is used with polling methods to receive upstream data from the set top terminals. The invention initiates such upstream data retrieval, gathers all data received and compiles viewer demographics information and programs watched information. The invention processes this data and information to generate packages of advertisements, as well as account and billing reports, targeted towards each set top terminal. The invention uses upstream data reception hardware, databases and processing hardware and software to accomplish these functions.

NETWORK CONTROLLER FOR CABLE TELEVISION DELIVERY SYSTEMS

RELATED APPLICATIONS

This application is a divisional of Canadian application serial number 2,447,895, which is a divisional of application serial number 2,151,459 the Canadian National Phase of International patent application PCT/US93/11616 filed 2 December 1993 (02.12.93).

TECHNICAL FIELD

The invention relates to television entertainment systems for providing television programming to consumer homes. More particularly, the invention relates to a network controller that monitors, controls and manages a television program delivery network from a cable headend.

BACKGROUND OF THE INVENTION

Advances in television entertainment have been primarily driven by breakthroughs in technology. In 1939, advances on Vladmir Zworykin's picture tube provided the stimulus for NBC to begin its first regular broadcasts. In

1975, advances in satellite technology provided consumers with increased programming to homes.

Many of these technology breakthroughs have produced inconvenient systems for consumers. One example is the ubiquitous three remote control home, having a separate and unique remote control for the TV, cable box and VCR. More recently, technology has provided cable users in certain parts of the country with 100 channels of programming. This increased program capacity is beyond the ability of many consumers to use effectively. No method of managing the program choices has been provided to consumers.

Consumers are demanding that future advances in television entertainment, particularly programs and program choices, be presented to the consumer in a user friendly manner. Consumer preferences, instead of technological breakthroughs, will drive the television entertainment market for at least the next 20 years. As computer vendors have experienced a switch from marketing new technology in computer hardware to marketing better useability, interfaces and service, the television entertainment industry will also experience a switch from new technology driving the market to consumer useability driving the market.

Consumers want products incorporating new technology that are useful, and will no longer purchase new technology for the sake of novelty or status. Technological advances in sophisticated hardware are beginning to surpass the capability of the average consumer to use the new technology. Careful engineering must be done to make entertainment products incorporating new technology useful and desired by consumers.

In order for new television entertainment products to be successful, the products must satisfy consumer demands.

TV consumers wish to go from limited viewing choices to a variety of choices, from no control of programming to complete control. Consumers wish to advance from cumbersome and inconvenient television to easy and convenient television and keep costs down. Consumers do not wish to pay for one hundred channels when due to lack of programming information, they seldom, if ever, watch programming on many of these channels. Viewers wish their programming to be customized and targeted to their needs and tastes.

The concepts of interactive television, high definition television and 300 channel cable systems in consumer homes will not sell if they are not packaged, delivered and presented in a useable fashion to consumers. Consumers are already being bombarded with programming options, numerous, "free" cable channels, subscription cable channels and pay-per-view choices. Any further increase in TV entertainment choices, without a user friendly presentation and approach, will likely bewilder viewers with a mind-numbing array of choices.

The TV industry has traditionally marketed and sold its programs to consumers in bulk, such as continuous feed broadcast and long-term subscriptions to movie channels. The TV industry is unable to sell its programming in large quantities on a unit per unit basis, such as the ordering of one program. Consumers prefer a unit sales approach because it keeps costs down and allows the consumer to be more selective in their viewing.

In today's television world, networks manage the program lineup for individual channels. Each network analyzes ratings for television shows and determines the appropriate schedule or program lineup to gain market share and revenue from advertising. Program ratings are determined using a test group of viewers and statistical analysis methods. Since each channel is in competition with every other channel, there is no coordinated effort to organize television programming in a manner that primarily suits the viewers.

Advertising has become equally annoying, with viewers being "forced" to watch television commercials for goods and services that are neither needed nor desired. As a result, consumers have become impatient and dissatisfied with today's television delivery systems. Equally problematic, these television delivery systems do not have the capabilities or features necessary to operate in the digital environment. Consequently, advances in digital technology call for a new television program delivery system that is capable of satisfying varying consumer and viewer needs.

Existing cable headends are unequipped for the transition to a digital system. These cable headends have no means for monitoring and controlling the large numbers of program signals and advertisements that will eventually be passed on to both consumers and viewers. These cable headends are unequipped to manage account and billing information for set top terminals without relying on telephone lines. In addition, these cable headends have no means for targeting advertisements to particular consumers and viewers.

What is needed is a network controller for a digital cable headend used in a television delivery system.

What is needed is a versatile network controller for a cable headend.

What is needed is a network controller for use in a cable headend that is capable of operating in both the digital and analog environment.

What is needed is certain components of a network controller for a digital cable headend used in a cable television delivery system.

What is needed is a network controller capable of controlling multiple video/audio program signals received by a cable headend from a satellite transponder.

What is needed is a network controller that can control the routing of both analog and digital video/audio program signals from cable headend to viewer homes.

What is needed is a network controller component for a cable headend that controls the combining the digital video/audio signals.

What is needed is a network controller that creates tiered programming by combining various digital video/audio signals.

What is needed is a network controller for a cable headend that accommodates different bandwidth availability between cable headend and certain viewer homes.

What is needed is a network controller capable of modifying program control information received from an external source.

What is needed is a network controller capable of targeting video to viewers.

What is needed is a network controller capable of targeting television commercials to specific consumers and viewers.

What is needed is a network controller capable of gathering information on programs watched by viewers.

What is needed is a better method of determining program ratings.

What is needed is a network controller capable of managing account and billing information.

The present invention is addressed to fulfill these needs.

SUMMARY OF INVENTION

The present invention is a network controller for a television delivery system. The network controller is the central component that provides monitoring and control of set top terminals in a television delivery system. The network controller is a key component of a digital cable television delivery system. The network controller of the present invention provides much greater capability and flexibility than existing cable headend control equipment.

The network controller of the preferred embodiment performs all its cable network monitoring and control of set top terminals within the cable headend. The cable headend receives and processes digitally compressed program signals before the signals are relayed to each set top terminal. Each cable headend site is equipped with multiple satellite receiver dishes and a signal processor.

As an intermediary between the set top terminals and the program delivery system's operations center (or other remote site), the cable headend relies on the network controller to perform key cable system operations. In particular, the network controller accommodates regional programming needs by working with other cable headend components. The network controller also performs the system control functions for the cable system.

The primary function of the network controller is to manage the configuration of set top terminals and process signals received from the set top terminals. In the preferred embodiment, the network controller monitors, among other things, automatic poll-back responses from the set top terminals remotely located at each subscribers' home. The polling and automatic report-back cycle occurs frequently enough to allow the network controller to maintain accurate account and billing information as well as monitor authorized channel access.

In the simplest embodiment, information to be sent to the network controller will be stored in RAM within each subscriber's set top terminal and will be retrieved only upon polling by the network controller. Retrieval may, for example, occur on a daily, weekly or monthly basis. The network controller allows the system to maintain complete information on all programs watched using a particular set top terminal.

The network controller is also able to respond to the immediate needs of a set top terminal, or a group of set top terminals. The network controller can modify a program signal received from the program delivery system's operations center before the program signal is transmitted to the set top terminal. Therefore, the network controller enables the delivery system to adapt to the specific requirements of individual set top terminals when information on these requirements cannot be provided to the operations center in advance. In other words, the network controller is able to perform "on the fly programming" changes. With this capability, the network controller can handle sophisticated local programming needs such as interactive television services, split screen video, and selection of different foreign languages for the same video. In addition, the network controller controls and monitors all compressors and decompressors in the system.

The network controller makes use of a number of software routines that assist the network controller to perform its major functions. One of the major routines assists the network controller to modify the program control information so that changes and additions in programming and advertisements can be accommodated. Such changes and additions include set top terminal access authorizations and deauthorizations.

A set top terminal data gathering routine allows the network controller to schedule and perform polling of all set top terminals operating in the system. The software also provides the network controller with a means of processing status reports received from set top terminals in response to polling requests.

A video targeting routine 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 generates packages of advertisements targeted towards each viewer.

Finally, an additional routine correlates the programs accessed with pricing information to generate billing reports that can be sent to a given set top terminal over the cable distribution network. Aside from this routine, the network controller accommodates other methods of billing and account maintenance, such as through the use of remote billing sites.

The present invention is not only able to operate in the digital environment but also introduces many new features to television program delivery and cable headend control.

In one aspect, the present invention provides an apparatus for targeting advertising to at least one subscriber, comprising by: means for gathering programs watched data from a subscriber; and a processor operably connected to the gathering means and including: means for analyzing the gathered programs watched data by determining the frequency of programs watched by the subscriber or the viewing habits of the subscriber; means for correlating the analyzed programs watched data with at least one advertisement by comparing the programs watched data with advertisements; means for selecting at least one correlated advertisement.

In a further aspect, the present invention provides an apparatus for providing digital program signals to subscriber locations, capable of inserting local availability signals and selecting digital program signals received from outside sources, comprising: a means for receiving digital program signals; digital logic circuitry, connected to the receiving means, wherein digital program signals may be inserted and wherein digital programs may be selected; a processor, operably connected to the digital logic circuitry, wherein insertion of local availability signals are controlled and wherein the selection of digital program signals are controlled; and means for sending programs, operably connected to the digital logic circuitry, wherein digital program signals are sent to subscriber locations, and wherein digital program signals that have been inserted or selected are sent to subscriber locations.

In a still further aspect, the present invention provides an apparatus for use in a program delivery system having a plurality of program signals, wherein the apparatus receives directions for local insertion from a remote source, comprising: a receiver, wherein directions for local insertion from a remote source are received; a database, containing information on a plurality of programs; a processor, operably connected to the receiver and database, wherein the directions from the remote source are processed to generate a processed signal; a local inserter, connected to the processor, wherein local programming is inserted into a digital program signal using the processed signal from the processor; and a modulator, wherein a plurality of program signals including the program signal carrying the inserted local programming is modulated.

In a further aspect, the present invention provides an apparatus that gathers information related to television viewing habits, comprising: a first processor that provides information related to which television programming was viewed; a first memory coupled to the first processor that stores the information; a second processor that receives the information from the first memory and arranges the information as programs watched data; and a second memory coupled to the second processor that stores the programs watched data.

In a further aspect, the present invention provides an apparatus that gathers programs watched information from set top terminals in a broadcast television program delivery system, comprising: means for gathering the programs watched information; means for storing the gathered programs watched information; means for sorting the stored programs watched information by a plurality of program categories; and means for ranking the plurality of program categories by frequency of programs watched in each category, wherein program categories with more programs watched are ranked higher than program categories with fewer programs watched.

In a still further aspect, the present invention provides a method for gathering programs watched information, comprising: gathering programs watched data for set top terminals in a broadcast television program delivery system; storing the gathered programs watched data in a database; accessing the stored programs watched data; and counting the accessed programs watched data to determine the frequency of programs watched by the set top terminals, wherein programs watched counts are arranged in a programs watched matrix.

In a further aspect, the present invention provides a television system that receives and stores data related to television viewing habits at a set top terminal, comprising; a receiver operably connected to the set top terminal, the receiver receiving programs watched information from the set top terminal; a processor coupled to the receiver, the processor processing the programs watched information to generate a programs watched matrix; and a database coupled to the processor that stores the programs watched matrix, wherein the processor is capable of intelligent selection of programs to be sent to the set top terminal.

In a further aspect, the present invention provides a network controller for use with digital signal processing equipment in a cable headend capable of remotely monitoring and controlling a plurality of set top terminals in a cable television program delivery system using network control data, each set top terminal receiving a plurality of information fields in a control information stream, the control information stream being produced using a program control information signal received from a remotely located source, the network controller comprising: an interface means, connected to the digital signal processing equipment, for receiving and transferring control information, wherein the program control information signal is received from the digital signal processing equipment on a scheduled basis and the control information stream produced using the program control information signal is transferred to the digital signal processing equipment for distribution over the cable television system; a means for storing the network control data, wherein the stored network control data includes data on television programs; a means for accessing the stored network control data; and a means for generating the control information stream using the received program control information signal and the accessed network control data, wherein the received program information signal carries data on packaged programs or menu content, whereby the information fields of the control information stream are formed by modifying the data on packaged programs or menu content and whereby the control information stream can subsequently be distributed to multiple set top terminals.

In a still further aspect, the present invention provides a method for using digital signal processing equipment in a cable headend to remotely monitor and control a plurality of set top terminals in a cable television system, each set top terminal adapted to transmit a set top terminal status report and to receive a control information stream produced from a program control information signal that communicates data in an information field, the method comprising the steps of: interfacing with the digital signal processing equipment, wherein the program control information signal is received from the digital signal processing equipment on a scheduled basis and the control information stream produced using the program control information signal is transferred back to the signal processing equipment for distribution over the cable television system; receiving set top terminal status reports, wherein the received set top terminal status reports provide the set top terminals with an upstream data transmission capability; storing network control data, wherein the stored network control data includes data on programs and contents of the received set top terminal status reports; accessing the stored network control data; and generating the control information stream using the received program control information signal and the accessed network control data, wherein the information fields of the control information stream are created and the control information stream is subsequently distributed to multiple set top terminals.

In a still further aspect, the present invention provides a network controller for use with digital signal processing equipment in a cable headend capable of monitoring and controlling a plurality of set top terminals in a cable television program delivery system that uses control information and network control data, wherein entries may be made and viewed by an operator, comprising; an interface, connected to the digital signal processing equipment, wherein the interface receives and transfers control information to the digital signal processing equipment; memory, wherein the memory stores the network control data including data on television programs in a database format and wherein the stored network control data may be accessed; a network controller CPU, connected to the memory and the interface, for generating a control information stream using the stored network control data from the memory, whereby information fields for the control information stream are formed and the control information stream is subsequently distributed to set top terminals; input device, connected to the network controller CPU, wherein entries may be made; and a display, connected to the network controller CPU, for viewing data and entries made on the input device.

In a further aspect, the present invention provides a network controller for remotely managing account and billing information over a cable television distribution network characterized by: a means for creating a polling request message; a means for transmitting the polling request message to set top terminals, wherein the polling request message directs each set top terminal to initiate upstream data transmission of set top terminal status reports over the cable distribution network; a means for receiving the set top terminal status reports, wherein the set top terminal status reports include programs watched data; a means for storing network control data, wherein the stored network control data includes programs watched data and price category data; a means for correlating the programs watched data with the price category data to produce price correlations; and a means for generating at least one billing report based on the produced price correlations; and a means for transmitting the billing reports to the set top terminals for display on a television screen.

In a still further aspect, the present invention provides anetwork controller for use with a cable headend capable of remotely targeting categories of advertisements to groups of set top terminals in a cable television system, the network controller characterized by: a means for gathering programs watched data for each set top terminal in the cable television system; a means for counting the gathered programs watched data to determine the frequency of programs watched by each set top terminal in the cable television program delivery system, wherein the programs watched counts are compiled by program category and time slot; a means for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each category advertisements includes one or more advertisements available for targeting to the set top terminals; and a means for transmitting the set top terminal group information to the set top terminal in a control information stream that instructs the set top terminal in selecting targeted advertisements for display during viewing of programs.

In a still further aspect, the present invention provides a etwork controller for remotely monitoring a plurality of set top terminals in a cable television system, the network controller characterized by: a means for receiving a program control information signal from digital signal processing equipment, wherein the digital signal processing equipment is co-located with the receiving means at the cable headend and the program control information signal is received by the digital signal processing equipment from a remotely located source; a means for storing network control data, wherein the stored network control data includes data on television programs; a means for accessing the stored network control data; a means for generating a control information stream using the received program control information signal and the accessed network control data, wherein the received program information signal carries data on packaged programs or menu content and whereby information fields of the control information stream are created by modifying the data on packaged programs or menu content based on the stored network control data so that the created information fields contain data on modified packages of programs or modified menu content, including program category and menu assignment information; a means for transferring the control information stream to the digital signal processing equipment, whereby the control information stream is subsequently distributed by the digital signal processing equipment to multiple set top terminals; a means for gathering programs watched data for each set top terminal in the cable television program delivery system, wherein the gathered programs watched data is stored in the storing means so that it can be accessed by the accessing means; a means for counting the accessed programs watched data to determine the frequency of programs watched by each set top terminal in the cable television program delivery system, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot; and a means, connected to the generating means, for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each advertisements category includes advertisements available for targeting to the set top terminals, and wherein the created set top terminal group information is included as data in the control information stream.

In a still further aspect, the present invention provides an apparatus for a program signal delivery system wherein digital program signals are used comprising: a receiver, wherein program signals are received; a demultiplexor, connected to the receiver, wherein program signals are demultiplexed for processing; digital logic, connected to the demultiplexer, wherein demultiplexed signals are processed; a control processing unit, operably connected to the digital logic, wherein the processing of the demultiplexed signals in the digital logic is controlled; multiplexor, connected to the digital logic, wherein processed signals are multiplexed; and a modulator, connected to the multiplexer, wherein multiplexed signals are modulated and wherein the modulated signal is forwarded to subscribers.

In a still further aspect, the present invention provides an apparatus for use with a network having subscribers that uses two-way communications allowing upstream data transmission from the subscriber comprising: a two-way communication system comprising one or more diplex filters, wherein the diplex filters facilitate upstream data transmissions; a signal processor, operably connected to the two-way communication system, wherein signals are processed before being sent downstream to subscribers; an RF combiner, connected to the two-way communication system, wherein upstream data transmissions are received from the two-way communication system; a network controller, connected to the RF combiner, wherein upstream data transmissions are received from the RF combiner; and wherein program control information is used, and the network controller comprises a processing unit, wherein program control information is processed.

In a still further aspect, the present invention provides a computer system for a network which controls video and audio delivery to subscriber locations, comprising: a network connecting subscriber locations; a plurality of terminals, connected to the network, to allow subscribers at subscriber locations to access the network; a processor, connected to the network and remotely located from the subscriber locations, wherein the processor generates control signals to control the distribution of video and audio signals to the plurality of terminals located at subscriber locations; and signal processing equipment, which communicates with the processor through a non-network connection, adapted to output audio and video signals to a plurality of terminals.

In a still further aspect, the present invention provides a network controller for use in a cable system for choosing text overlays, wherein a subscriber may request a text overlay representing the audio associated with the video to a program, wherein a set top terminal transmits the subscriber requests comprising: a means for receiving subscriber requests; a processor to identify the subscriber request for text overlays from a plurality of subscriber requests; a means for choosing the corresponding text overlay from a plurality of text overlays; and a means for transmitting text overlay to the set top terminal for display.

In a still further aspect, the present invention provides a system for choosing audio channels, wherein more than one audio channel may exist for a video, comprising: a signal processor, wherein a plurality of digital video signals and audio channels are received, and wherein a received video signal may have multiple audio channels; a controller, wherein the controller selects one or more received audio channels, wherein one or more audio channels that correspond to a received video are selected; a distribution system, operably connected to the controller, wherein the selected audio channels are distributed and a listener receives one or more audio channels; and wherein a program control information signal is used by the controller to select one or more received audio channels.

In a still further aspect, the present invention provides an apparatus that gathers programs watched data, comprising: a plurality of terminals connected to corresponding televisions and to a television program delivery system, each of the terminals including a memory that stores program access information; and a receiver coupled to the plurality of terminals, the receiver receiving the program access information, wherein the program access information is stored as programs watched data.

In a still further aspect, the present invention provides a system that gathers programs watched data in a broadcast television delivery system, comprising: means for gathering programs watched data from one or more set top terminals in a broadcast television delivery system; a databases that stores the gathered programs watched data; means for accessing the stored programs watched data; means for counting the accessed programs watched data to determine programs watched counts corresponding to the frequency of programs watched by the one or more set top terminals in the broadcast television delivery system, wherein the programs watched counts are arranged in at least one programs watched matrix; means for creating set top terminal group information indicating a group assignment for a set top terminal by correlating the programs watched counts with categories of videos, wherein the video categories include videos available for sending to the set top terminal; and means for transmitting the set top terminal group information to the set top terminal in a control information stream that instructs the set top terminal in selecting the videos for display.

In a further aspect, the present invention provides a method of gathering programs watched information from a set top terminal that acquires programs from a programming source, comprising: receiving a plurality of programs at the set top terminal; selecting a program at the set top terminal from the plurality of programs; monitoring the program selected at the set top terminal; generating programs watched information relating to the selected program; and storing the generated programs watched information.

In a still further aspect, the present invention provides an apparatus for data processing related to program advertisements wherein data on viewer profiles and programs watched data is used, comprising: 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, and wherein the information on advertisements identifies individual advertisements by name; a secondary memory location, operably connected to the receiver, wherein data on viewer profiles is stored; 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; and a third memory location, operably connected to the central processing unit, wherein programs watched data is stored, wherein the programs watch data includes data related to all programs watched at a set top terminal, and wherein the central processing unit processes the information on advertisements, viewer profile data and programs watched data to schedule advertisements.

In a further aspect, the present invention provides a method for data processing related to program advertisements wherein data on viewer profiles and programs watched data is used, comprising: receiving viewer profile data at a receiver from one or more remote locations; storing information on advertisements in a first memory location operably connected to the receiver, wherein the information on advertisements identifies individual advertisements by name; storing data on viewer profiles in a secondary memory location operably connected to the receiver; processing the information on advertisements and viewer profile data, and scheduling at least one advertisement, with a central processing unit operably connected to the first and second memory locations, wherein at least one advertisement is scheduled to be inserted into programming; storing programs watched data in a third memory location operably connected to the central processing unit, wherein the programs watch data includes data related to all programs watched at a set top terminal; and processing the information on advertisements, viewer profile data and programs watched data, using the central processing unit, to schedule advertisements.

In a still further aspect, the present invention provides a method for targeting advertising to at least one subscriber, the method comprising 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.

In a further aspect, the present invention provides a system for targeting advertising to at least one subscriber comprising: a means for gathering programs watched data from a subscriber; a processor, operably connected to the gathering means, comprising 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.

In a still further aspect, the present invention provides a system for targeting advertising comprising: a means for gathering programs watched data from subscribers to be targeted; a processor comprising: 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.

In a further aspect, the present invention provides a network controller for use in a cable television program delivery system for targeting advertising, the network controller comprising: a means for gathering programs watched data from set top terminals; a means for storing the gathered programs watched data; a means for accessing the stored programs watched data; a means for counting the accessed programs watched data to determine the frequency of programs watched by the set top terminals, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot; a means for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each advertisement category includes at least one advertisement; and a transmitter, wherein the transmitter transmits the set top terminal group information to the set top terminals in a control information stream that instructs the set top terminals in selecting advertisements for display during viewing of programs.

In a still further aspect, the present invention provides a method of targeting advertisements comprising 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.

In a further aspect, the present invention provides an apparatus for targeting advertising to a subscriber, the apparatus comprising: 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.

In a still further aspect, the present invention provides a method of targeting programs comprising 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.

In a further aspect, the present invention provides an apparatus for targeting video, wherein data on programs watched is used to determine a subscriber's preference, comprising: 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.

In a still further aspect, the present invention provides a method for targeting advertising using subscriber information, comprising: gathering programs watched data from a subscriber; analyzing the programs watched data to determine a 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 a first advertisement and a second advertisement from the correlated advertising categories; transmitting the first advertisement for display to a first subscriber and the second advertisement for display to a second subscriber; and gathering advertisements watched data from the first and the second subscribers.

In a further aspect, the present invention provides a method for targeting advertisements to television viewers based on television viewer information, comprising: receiving gathered television viewer information; analyzing the received viewer information; correlating the analyzed viewer information to advertisement categories; selecting a first advertisement from a first advertisement category; selecting one or more subsequent advertisements from one of the first advertisement category and one or more subsequent advertisement categories; providing the selected first advertisement and the one or more subsequent advertisements to television terminals; providing a switching plan that directs the television terminals to display one of the provided first advertisement and the one or more subsequent advertisements during an advertisement spot in a television program.

In a still further aspect, the present invention provides a method for targeting advertisements to television viewers in a television program delivery system, comprising: determining advertisement categories; mapping television terminals to the determined advertisement categories, wherein one or more of the television terminals are assigned to groups for one or more of the advertisement categories; and mapping television programs to the groups.

Further aspects of the invention will become apparent upon reading the following detailed description and drawings, which illustrate the invention and preferred embodiments of the invention.

It is an object of this invention to provide a network controller for a television delivery system.

It is an object of this invention to provide a network controller for a digital cable headend used in a cable television delivery system.

It is an object of this invention to provide certain needed components of a network controller for a digital cable headend used in a cable television delivery system.

It is an object of this invention to provide a versatile network controller for a cable headend.

It is an object of this invention to provide a network controller for use in a cable headend that is capable of operating in both the digital and analog environment.

It is an object of this invention to provide a network controller capable of controlling multiple video/audio program signals received by a cable headend.

It is an object of this invention to provide a network controller that can control the routing of both analog and digital video/audio program signals from cable headend to viewer homes.

It is an object of this invention to provide a network controller component for a cable headend that controls the combining the digital video/audio signals.

It is an object of this invention to provide a network controller that creates tiered programming by combining various digital video/audio signals.

It is an object of this invention to provide a network controller for a cable headend that accommodates different bandwidth availability between cable headend and certain viewer homes.

It is an object of the invention to provide a network controller capable of modifying program control information received from an external source.

It is an object of the invention to provide a network controller capable of targeting specific video/audio to specific viewers.

It is an object of the invention to provide a network controller capable of targeting television commercials to specific consumers and viewers.

It is an object of the invention to provide a network controller capable of retrieving data gathered at set top terminals.

It is an object of this invention to provide a network controller capable of managing account and billing information.

These and other objects and advantages of the invention will become obvious to those skilled in the art upon review of the following description, the attached drawings and appended claims.

DESCRIPTION OF THE DRAWINGS

Figure 1 is a diagram of the primary components of the television delivery system.

Figure 2 is an overview of the television delivery system operations.

Figure 3 is a schematic of the operation of the primary components of the system.

Figure 4 is a diagram of the primary components of the cable headend.

Figure 5 is a diagram of the cable headend showing the primary components of the network controller.

Figure 6a is a schematic of a basic cable headend having network controller components.

Figure 6b is a schematic of an alternative embodiment of Figure 6a.

Figure 7 is a detailed diagram of the components of the cable headend.

Figure 8a is a drawing of a broadcast television menu screen to be displayed on a set top terminal.

Figure 8b is a drawing of a hit movie menu screen to be displayed on a set top terminal.

Figure 8c is a drawing of a hit movie description menu screen to be displayed on a set top terminal.

Figure 9a is a diagram for out-of-band two-way data transmission for a digital/analog headend.

Figure 9b is a diagram for in-band two-way data transmission for a digital/analog headend.

Figure 10a is a diagram of the polling request message format.

Figure 10b is a diagram of the polling response message format with an expanded view of the programs accessed block field.

Figure 11 is a diagram of the network controller CPU and its relational components.

Figure 12 is diagram of the network control database structure.

Figure 13 is a diagram of the relationship between the major software routines.

Figure 14 is a block diagram of the software flow chart for the Modifying PCI routine.

Figure 15 is a block diagram of the software flow chart for the Polling Cycle routine.

Figure 16 is a diagram of a sample programs watched matrix.

Figure 17 is a block diagram of the software flow chart for the Basic Advertisement Targeting routine.

Figure 18 is a block diagram of the subroutine flow chart for processing programs watched matrices through correlation algorithms.

Figure 19 is a diagram of the subroutine flow chart for determining final groupings of set top terminals.

Figure 20a is a diagram showing a sample assignment of advertising channels to set top terminal groups watching particular categories of programs.

Figure 20b is a diagram assigning available bandwidth for multiple advertising channels.

Figure 21 is a diagram of the software flow chart for an alternative to the Basic Advertisement Targeting routine.

Figure 22 is a diagram of the software flow chart for the Account/Billing routine.

Figure 23 is a diagram of an embodiment that uses remote statistical and billing sites.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

A. Television Program Delivery System Description

1. Introduction

Figure 1 shows the present invention as part of an expanded cable television program delivery system 200 that dramatically increases programming capacity using compressed transmission of television program signals. Developments in digital bandwidth compression technology now allow much greater throughput of television program signals over existing or slightly modified transmission media.

The program delivery system 200 shown provides subscribers with a user friendly interface to operate and exploit a six-fold or more increase in current program delivery capability.

Subscribers are able to access an expanded television program package and view selected programs through a menu-driven access scheme that allows each subscriber to select individual programs by sequencing a series of menus. The menus are sequenced by the subscriber using simple alpha-numeric and iconic character access or moving a cursor or highlight bar on the TV screen to access desired programs by simply pressing a single button, rather than recalling from memory and pressing the actual two or more digit numeric number assigned to a selection. Thus, with the press of a single button, the subscriber can advance from one menu to the next. In this fashion, the subscriber can sequence the menus and select a program from any given menu. The programs are grouped by category so that similar program offerings are found on the same menu.

2. Major System Components

In its most basic form, the system uses a program delivery system 200 in conjunction with a conventional concatenated cable television system 210. The program delivery system 200 generally includes (i) at least one operations center 202, where program packaging and control information are created and then assembled in the form of digital data, (ii) a digital compression system, where the digital data is compressed, combined/multiplexed, encoded, and mapped into digital signals for satellite transmission to the cable headend 208, and (iii) a set of in-home decompressors. The program delivery system 200 transports the digital signals to the cable headend 208 where the signals are transmitted through a concatenated cable television system 210. Within the cable headend 208, the received signals may be decoded, demultiplexed, managed by a local central distribution and switching mechanism, combined and then transmitted to the set top terminal 220 located in each subscriber's home over the cable system 210. Although concatenated cable systems 210 are the most prevalent transmission media to the home, telephone lines, cellular networks, fiberoptics, Personal Communication Networks and similar technology for transmitting to the home can be used interchangeably with this program delivery system 200.

The delivery system 200 has a reception region 207 with an in-home decompression capability. This capability is performed by a decompressor housed within a set top terminal 220 in each subscriber's home. The decompressor remains transparent from the subscriber's point of view and allows any of the compressed signals to be demultiplexed and individually extracted from the composite data stream and then individually decompressed upon selection by the subscriber. The decompressed video signals are converted into analog signals for television display. Such analog signals include NTSC formatted signals for use by a standard television. Control signals are likewise extracted and decompressed and then either executed immediately or placed in local storage such as a RAM. Multiple sets of decompression hardware may be used to decompress video and control signals. The set top terminal 220 may then overlay or combine different signals to form the desired display on the subscriber's television. Graphics on video or picture-on-picture are examples of such a display.

Although a single digital compression standard (e.g., MPEG) may be used for both the program delivery system 200 and the concatenated cable system 210, the compression technique used may differ between the two systems. When the compression standards differ between the two media, the signals received by the cable headend 208 must be decompressed before transmission from the headend 208 to the set top terminals 220. Subsequently, the cable headend 208 must recompress and transmit the signals to the set top terminal 220, which would then decompress the signals using a specific decompression algorithm.

The video signals and program control signals received by the set top terminal 220 correspond to specific television programs and menu selections that each subscriber may access through a subscriber interface. The subscriber interface is a device with buttons located on the set top terminal 220 or on a portable remote control 900. In the preferred system embodiment, the subscriber interface is a combined alpha-character, numeric and iconic remote control device 900, which provides direct or menu-driven program access. The preferred subscriber interface also contains cursor movement and go buttons as well as alpha, numeric and iconic buttons. This subscriber interface and menu arrangement enables the subscriber to sequence through menus by choosing from among several menu options that are displayed on the television screen. In addition, a user may bypass several menu screens and immediately choose a program by selecting the appropriate alpha-character, numeric or iconic combinations on the subscriber interface. In the preferred embodiment, the set top terminal 220 generates the menus that are displayed on the television by creating arrays of particular menu templates, and the set top terminal 220 displays a specific menu or submenu option for each available video signal.

3. Operations Center and Digital Compression System

The operations center 202 performs two primary services, packaging television programs and generating the program control information signal. At the operations center 202, television programs are received from external program sources in both analog and digital form. Figure 2 shows an embodiment of the operations center receiving signals from various external sources 212. Examples of the external program sources are sporting events, children's programs, specialty channels, news or any other program source that can provide audio or visual signals. Once the programs are received from the external program sources, the operations center 202 digitizes (and preferably compresses) any program signals received in analog form. The operations center 202 may also maintain an internal storage of programs. The internally stored programs may be in analog or digital form and stored on permanent or volatile memory sources, including magnetic tape or RAM. Subsequent to receiving programming, the operations center 202 packages the programs into the groups and categories which provide the optimal marketing of the programs to subscribers. For example, the operations center 202 may package the same programs into different categories and menus for weekday, prime-time viewing and Saturday afternoon viewing. Also, the operations center 202 packages the television programs in a manner that enables both the various menus to easily represent the programs and the subscribers to easily access the programs through the menus.

The packaging of the digital signals is typically performed at the operations center 202 by computer assisted packaging equipment (CAP). The CAP system normally includes at least one computer monitor, keyboard, mouse, and standard video editing equipment. A programmer packages the signals by entering certain information into the CAP. This information includes the date, time slot, and program category of the various programs. The programmer and the CAP utilize demographic data and ratings in performing the packaging tasks. After the programmer selects the various programs from a pool of available programs and inputs the requisite information, the programmer, with assistance from the CAP, can select the price and allocate transponder space for the various programs. After the process is complete, the CAP displays draft menus or program schedules that correspond to the entries of the programmer. The CAP may also graphically display allocation of transponder space. The programmer may edit the menus and transponder allocation several times until satisfied with the programming schedule. During the editing, the programmer may direct the exact location of any program name on a menu with simple commands to the CAP.

The packaging process also accounts for any groupings by satellite transponder which are necessary. The operations center 202 may send different groups of programs to different cable headends 208 and/or set top terminals 220. One way the operations center 202 may accomplish this task is to send different program packages to each transponder. Each transponder, or set of transponders, then relays a specific program package to specific cable headends 208 and/or set top terminals 220. The allocation of transponder space is an important task performed by the operations center 202.

The operations center 202 may also "insert" directions for filling local available program time in the packaged signal to enable local cable and television companies to fill the program time with local advertising and/or local programming. Consequently, the local cable headends 208 are not constrained to show only programs transmitted from the operations center 202. New set top converters will incorporate both digital and analog channels. Therefore, the cable headend 208 may combine analog signals with the digital signals prior to transmitting the program signals to the set top terminals 220.

After the CAP packages the programs, it creates a program control information signal to be delivered with the program package to the cable headend 208 and/or set top terminal 220. The program control information signal contains a description of the contents of the program package, commands to be sent to the cable headend 208 and/or set top terminal 220, and other information relevant to the signal transmission.

In addition to packaging the signal, the operations center 202 employs digital compression techniques to increase existing satellite transponder capacity by at least a 4:1 ratio, resulting in a four-fold increase in program delivery capability. A number of digital compression algorithms currently exist which can achieve the resultant increase in capacity and improved signal quality desired for the system. The algorithms generally use one or more of three basic digital compression techniques: (1) within-frame (intraframe) compression, (2) frame-to-frame (interframe) compression, and (3) within carrier compression. Specifically, in the preferred embodiment, the MPEG 2 compression method is used. After digital compression, the signals are combined (multiplexed) and encoded. The combined signal is subsequently transmitted to various uplink sites 204.

There may be a single uplink site 204 or multiple uplink sites (represented by 204', shown in phantom in Figure 1) for each operation center 202. The uplink sites 204 may either be located in the same geographical place or may be located remotely from the operations center 202. Once the composite signal is transmitted to the uplink sites 204, the signal may be multiplexed with other signals, modulated, upconverted and amplified for transmission over satellite. Multiple cable headends 208 may receive such transmissions.

In addition to multiple uplinks, the delivery system 200 may also contain multiple operations centers. The preferred method for using multiple operations centers is to designate one of the operations centers as a master operations center and to designate the remaining operations centers as slave operations centers. In this configuration, the master operations center coordinates various functions among the slave operations centers such as synchronization of simultaneous transmissions and distributes the operations workload efficiently.

4. <u>Cable Headend</u>

After the operations center 202 has compressed and encoded the program signals and transmitted the signals to the satellite, the cable headend 208 receives and further processes the signals before they are relayed to each set top terminal 220. Each cable headend site is generally equipped with multiple satellite receiver dishes. Each dish is capable of handling multiple transponder signals from a single satellite and sometimes from multiple satellites.

As an intermediary between the set top terminals 220 and the operations center 202 (or other remote site), the cable headend 208 performs two primary functions. First, the cable headend 208 acts as a distribution center, or signal processor, by relaying the program signal to the set top terminal 220 in each subscriber's home. In addition, the cable headend 208 acts as a network controller 214 by receiving information from each set top terminal 220 and passing such information on to an information gathering site such as the operations center 202.

Figure 3 shows an embodiment where the cable headend 208 and the subscriber's home are linked by certain communications media 216. In this particular embodiment, analog signals, digitally compressed signals, other digital signals and up-stream/interactivity signals are sent and received over the media 216. The cable headend 208 provides such signaling capabilities in its dual roles as a signal processor 209 and network controller 214.

As a signal processor 209, the cable headend 208 prepares the program signals that are received by the cable headend 208 for transmission to each set top terminal 220. In the preferred system, the signal processor 209 re-routes or demultiplexes and recombines the signals and digital information received from the operations center 202 and allocates different portions of the signal to different frequency ranges. Cable headends 208 which offer different subscribers different program offerings may allocate the program signals from the operations center 202 in various manners to accommodate different viewers. The signal processor 209 may also incorporate local programming and/or local advertisements into the program signal 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 operations center program signals. If the local cable system uses a compression 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. This process becomes less important as standards develop (i.e., MPEG 2). In addition, the signal processor 209 performs any necessary signal decryption and/or encryption.

As a network controller 214, the cable headend 208 performs the system control functions for the system. The primary function of the network controller 214 is to manage the configuration of the set top terminals 220 and process signals received from the set top terminals 220. In the preferred embodiment, the network controller 214 monitors, among other things, automatic poll-back responses from the set top terminals 220 remotely located at each subscribers' home. The polling and automatic report-back cycle occurs frequently enough to allow the network controller 214 to maintain accurate account and billing information as well as monitor authorized channel access. In the simplest embodiment, information to be sent to the network controller 214 will be stored in RAM within each subscriber's set top terminal 220 and will be retrieved only upon polling by the network controller 214. Retrieval may, for example, occur on a daily, weekly or monthly basis. The network controller 214 allows the system to maintain complete information on all programs watched using a particular set top terminal 220.

The network controller 214 is also able to respond to the immediate needs of a set top terminal 220 by modifying a program control information signal received from the operations center 202. Therefore, the network controller

214 enables the delivery system to adapt to the specific requirements of individual set top terminals 220 when the requirements cannot be provided to the operations center 202 in advance. In other words, the network controller 214 is able to perform "on the fly programming" changes. With this capability, the network controller 214 can handle sophisticated local programming needs such as, for example, interactive television services, split screen video, and selection of different foreign languages for the same video. In addition, the network controller 214 controls and monitors all compressors and decompressors in the system.

The delivery system 200 and digital compression of the preferred embodiment provides a one-way path from the operations center 202 to the cable headend 208. Status and billing information is sent from the set top terminal 220 to the network controller 214 at the cable headend 208 and not directly to the operations center 202. Thus, program monitoring and selection control will take place only at the cable headend 208 by the local cable company and its decentralized network controllers 214 (i.e., decentralized relative to the operations center 202, which is central to the program delivery system 200). The local cable company will in turn be in communication with the operations center 202 or a regional control center (not shown) which accumulates return data from the set top terminal 220 for statistical or billing purposes. In alternative system embodiments, the operations center 202 and the statistical and billing sites are collocated. Further, telephone lines with modems are used to transfer information from the set top terminal 220 to the statistical and billing sites.

5. Set Top Terminal

The set top terminal 220 is the portion of the delivery system 200 that resides in the home of a subscriber. The set top terminal 220 is usually located above or below the subscriber's television, but it may be placed anywhere in or near the subscriber's home as long as it is within the range of the subscriber's remote control device 900. In some aspects, the set top terminal 220 may resemble converter boxes already used by many cable systems. For instance, each set top terminal 220 may include a variety of error detection, decryption, and coding techniques such as anti-taping encoding. However, it will become apparent from the discussion below that the set top terminal 220 is able to perform many functions that an ordinary converter box cannot perform.

The set top terminal 220 has a plurality of input and output ports to enable it to communicate with other local and remote devices. The set top terminal 220 has an input port that receives information from the cable headend 208. In addition, the unit has at least two output ports which provide communications from the set top terminal 220 to a television and a VCR. Certain menu selections may cause the set top terminal 220 to send control signals directly to the VCR to automatically program or operate the VCR. Also, the set top terminal 220 contains a phone jack which can be used for maintenance, trouble shooting, reprogramming and additional customer features. The set top terminal 220 may also contain stereo/audio output terminals and a satellite dish input port.

Functionally, the set top terminal 220 is the last component in the delivery system chain. The set top terminal 220 receives compressed program and control signals from the cable headend 208 (or, in some cases, directly from the operations center 202). After the set top terminal 220 receives the individually compressed program and control signals, the signals are demultiplexed, decompressed, converted to analog signals (if necessary) and either placed in local storage (from which the menu template may be created), executed immediately, or sent directly to the television screen.

After processing certain signals received from the cable headend 208, the set top terminal 220 is able to store menu templates for creating menus that are displayed on a subscriber's television by using an array of menu templates. Before a menu can be constructed, menu templates must be created and sent to the set top terminal 220 for storage. A microprocessor uses the control signals received from the operations center 202 or cable headend 208 to generate the menu templates for storage. Each menu template may be stored in volatile memory in the set top terminal 220. When the set top terminal receives template information it demultiplexes the program control signals received from the cable headend 208 into four primary parts: video, graphics, program logic and text. Each menu template represents a different portion of a whole menu, such as a menu background, television logo, cursor highlight overlay, or other miscellaneous components needed to build a menu. The menu templates may be deleted or altered using control signals received from the operations center 202 or cable headend 208.

Once the menu templates have been stored in memory, the set top terminal 220 can generate the appropriate menus. In the preferred embodiment, the basic menu format information is stored in memory located within the set top terminal 220 so that the microprocessor may locally access the information from the set top terminal instead of from an incoming signal. The microprocessor next generates the appropriate menus from the menu templates and the other menu information stored in memory. The set top terminal 220 then displays specific menus on the subscriber's television screen that correspond to the inputs the subscriber selects.

If the subscriber selects a specific program from a menu, the set top terminal 220 determines on which channel the program is being shown, demultiplexes and extracts the single channel transmitted from the cable headend 208. The set top terminal 220 then decompresses the channel and, if necessary, converts the program signal to an analog NTSC signal to enable the subscriber to view the selected program. The set top terminal 220 can be equipped to decompress more than one program signal, but this would unnecessarily add to the cost of the unit since a subscriber will generally only view one program at a time. However, two or three decompressors may be desirable to provide picture-on-picture capability, control signal decompression, enhanced channel switching or like features.

In addition to menu information, the set top terminal 220 may also store text transmitted from the cable headend 208 or the operations center 202. The text may inform the subscriber about upcoming events, billing and account status, new subscriptions, or other relevant information. The text will be stored in an appropriate memory location depending on the frequency and the duration of the use of the textual message.

Also, optional upgrades are available to enhance the performance of a subscriber's set top terminal 220. These upgrades may consist of a cartridge or computer card (not shown) that is inserted into an expansion slot in the set top terminal 220 or may consist of a feature offered by the cable headend 208 or operations center 202 to which the user may subscribe. Available upgrades may include on line data base services, interactive multi-media services, access to digital radio channels, and other services.

In the simplest embodiment, available converter boxes such as those manufactured by General Instruments or Scientific Atlanta, may be modified and upgraded to perform the functions of a set top terminal 220. The preferred upgrade is a circuit card with a microprocessor which is electronically connected to or inserted into the converter box.

6. Remote Control Device

The primary conduit for communication between the subscriber and the set top terminal 220 is through the subscriber interface, preferably a remote control device 900. Through this interface, the subscriber may select desired programming through the system's menu-driven scheme or by directly accessing a specific channel by entering the actual channel number. Using the interface, the subscriber can navigate through a series of informative program selection menus. By using menu-driven, iconic or alpha-character access, the subscriber can access desired programs by simply pressing a single button rather than recalling from memory and pressing the actual channel number to make a selection. The subscriber can access regular broadcast and basic cable television stations by using either the numeric keys on the remote control 900 (pressing the corresponding channel number), or one of the menu icon selection options.

In addition to enabling the subscriber to easily interact with the cable system 200, the physical characteristics of the subscriber interface 900 should also add to the user friendliness of the system. The remote control 900 should easily fit in the palm of the user's hand. The buttons of the preferred remote control 900 contain pictorial symbols that are easily identifiable by the subscriber. Also, buttons that perform similar functions may be color coordinated and consist of distinguishing textures to increase the user friendliness of the system.

7. Menu-Driven Program Selection

The menu-driven scheme provides the subscriber with one-step access to all major menus, ranging from hit movies to sport specials to specialty programs. From any of the major menus, the subscriber can in turn access submenus and minor menus by cursor or alpha-character access.

There are two different types of menus utilized by the preferred embodiment, the Program Selection menus and the During Program menus. The first series of menus, Program Selection menus, consists of an Introductory, a Home, Major menus, and Submenus. The second series of menus, During Program menus, consists of two primary types, Hidden menus and the Program Overlay menus.

Immediately after the subscriber turns on the set top terminal 220, the Introductory menu welcomes the subscriber to the system. The Introductory menu may display important announcements from the local cable franchise, advertisements from the cable provider, or other types of messages. In addition, the Introductory menu can inform the subscriber if the cable headend 208 has sent a personal message to the subscriber's particular set top terminal 220.

After the Introductory menu has been displayed the subscriber may advance to the next level of menus, namely the Home menu. In the preferred embodiment, after a certain period of time, the cable system will advance the subscriber by default to the Home menu. From the Home menu, the subscriber is able to access all of the programming options. The subscriber may either select a program directly by entering the appropriate channel number from the remote control 900, or the subscriber may sequence through incremental levels of menu options starting from the Home menu. The Home menu lists categories that correspond to the first level of menus called Major menus.

If the subscriber chooses to sequence through subsequent menus, the subscriber will be forwarded to the Major menu that corresponds to the chosen category from the Home menu. The Major menus further refine a subscriber's search and help guide the subscriber to the selection of his choice.

From the Major menus, the subscriber may access several submenus. From each submenu, the subscriber may access other submenus until the subscriber finds a desired television program. Similar to the Major menu, each successive level of Submenus further refines the subscriber's search. The system also enables the subscriber to skip certain menus or submenus and directly access a specific menu or television program by entering the appropriate commands on the remote control 900.

The During program menus (including Hidden Menus and Program Overlay Menus) are displayed by the set top terminal 220 only after the subscriber has selected a television program. In order to avoid disturbing the subscriber, the set top terminal 220 does not display the Hidden Menus until the subscriber selects the appropriate option to display a Hidden Menu. The Hidden Menus contain options that are relevant to the program selected by the viewer. For example, a Hidden Menu may contain options that enable a subscriber to enter an interactive mode or escape from the selected program.

Program Overlay Menus are similar to Hidden Menus because they occur during a program and are related to the program being viewed. However, the Program Overlay Menus are displayed concurrently with the program selected by the subscriber. Most Program Overlay Menus are small enough on the screen to allow the subscriber to continue viewing the selected program comfortably.

B. <u>Network Controller Description</u>

1. <u>Monitoring and Control of Set Top Terminals</u>

Figure 4 shows the network controller 214 of the present invention as part of a digital cable headend 208 operating in an expanded cable television program delivery system, indicated generally at 200. The network controller 214 monitors program selections at subscribers' homes, maintains accurate account and billing information and authorizes both subscriber channel access and particular set top terminals 220 to operate in the system.

The network controller 214 performs its monitoring and control capability by working with other system components housed, in part, within the cable headend 208. These cable headend components include a cable headend receiver 203 and a signal processor 209. As shown in the Figure 4, digital RF program signals 205 are received and processed for further distribution to a subscriber's home through a set top terminal 220. The program signals 205 are digitally compressed and multiplexed signals that may be processed at the cable headend 208 or simply passed through to the cable distribution network. In the embodiment shown in Figure 4, the program signals 205 are received by the cable headend receiver 203 and transmitted to the signal processor 209.

The signal processor 209 prepares the program signals 205 that are received by the cable headend 208 for transmission to each set top terminal 220. In the preferred system, the network controller 214 supervises and, in some cases, instructs the signal processor 209 in routing the signals to subscribers. In this way, the network controller 214 and signal processor 209 work with one another to perform basic control functions in the cable television system 200. Typically, this work is accomplished by the transfer of control information, represented at 211, between the network controller 214 and the signal processor 209.

Although it is preferred that the signal processor 209 and network controller 214 be co-located at the cable headend 208, the network controller 214 may be remotely located from the cable headend 208, as long as it remains in communication with the signal processor 209 in order to exchange control information 211.

In many instances, the program signals 205 received from the operations center 202 must be modified prior to being sent to the set top terminals 220. These modifications to the program control information 211 are made by the network controller 214 working in conjunction with the signal processor 209 to send a set top terminal control information stream (STTCIS). From the signal processor 209, the network controller 214 receives the program signals 205, which include cable franchise specific information added by the operations center 202. The network controller 214 modifies the program signals 205, if necessary, and communicates the new information back to the signal processor 209. The signal processor 209 then forwards the information to the set top terminal 220 in the form of the STTCIS, arrow 215. In most instances, the network controller 214 will modify the program signals 205 by adding additional information; however, the program signals 205 can be passed through the cable headend 208 to the set top terminal 220 without any modification.

The signal processor 209 and network controller 214 are both capable of handling the addition of simple local availabilities (e.g., local advertisements) into the signal sent to the set top terminal 220. The network controller 214 is also capable of handling more sophisticated local programming needs such as targeting video commercials, infomercials, interactive programming and certain data services. The network controller 214 receives all electronic signals sent by the set top terminal 220, including those sent in response to interactive service requests and some data service requests. The network controller 214 coordinates the necessary switching and access to allow the subscriber to enjoy these services.

The network controller 214 has the capability of performing "on the fly programming" changes, assisting in (i) masking portions of subscriber's television screens (split screen video), (ii) selecting different audio signals for the same video (foreign languages), and (iii) interactive features. In addition, the network controller can create programming changes. For last minute changes to programming (such as for a local emergency or important regional events), an operator using the network controller 214 can modify the program signals 209 "on the fly" and change menus available to the subscriber. This accommodates short notice changes to program packaging that cannot be handled by the operations center 202 in advance.

In order to accommodate split screen techniques for promo and demo video (which will be described later), undesired video portions of the television or menu screen may be masked. The network controller 214 can send the necessary control information to inform the set top terminal 220 to mask portions of a specific channel's video. For example, a video channel with a split screen showing four separate videos would require a three-fourths mask to focus the viewer on the featured video clip.

Tiered programming allows different users to view different video even though they are "tuned" to the same channel. For example, the network controller 214 may know the demographics of its subscribers through a database generated, in part, from prior subscriber choices, an interactive selection, or other means. Using the demographics information, the network controller 214 may target commercials to the correct audience by showing different commercials to subscriber's with different demographics. Information on programs watched may also be used to target commercials. Even though subscribers will believe they are "tuned" to one channel, they will be switched to a different channel for the tiered video and targeted commercial. Alternatively, individual subscribers may be offered a menu with the option of several commercials from which to choose.

To accommodate foreign speaking subscribers, multiple audio channels for television programming may be provided. The subscriber may be shown menus of programs available in the subscriber's native language. The function of choosing the correct audio to correspond to the selected language may be handled by either the set top terminal 220 or the network controller 214 depending upon the configuration. Local programming in several languages or additional audio channels for a foreign language translation of a popular television program may be provided by the network controller 214. Using a picture-on-picture feature, sign language may be similarly made available to certain set top terminals 220 for the deaf. The sign language video may be transmitted to the set top terminal 220 on a separate channel. Also, a text overlay for the deaf may be easily produced on the lower part of the screen. The control signals for producing the text overlay may be handled by the network controller 214.

In other embodiments, the network controller 214 can act as a central computer and provide intra-set top terminal interactive games, inter-set top terminal interactive games, computer bulletin board type services, message services (Electronic mail), etc. For example, a subscriber may play war games with six of his (anonymous) fellow subscribers each in their own home each operating a separate tank. The network controller 214 gathers the players using set top terminal 220 communications and acts as the referee. The network controller software "plays" the game and generates the video control signals to be transmitted to the set top terminals 220. From the video control signals, the set top terminal generates a view of the playing field and shows movement of the tanks. Using a similar method, a bulletin board or message system can be set up to discuss a particular program such as "Twin Peaks Whodunit" for enthusiasts with set top terminals 220.

2. <u>Monitoring and Control of Cable Headend Signal Processor</u>

Figure 5 shows the network controller's major components and how these components relate with other components of the cable system 200. The network controller's internal components include a network controller CPU 224, databases 226, control receiver 228, local memory 230 and telephone modem 232. The network controller's CPU 224 and databases 226 may be accessed through an operator control station, which may include peripherals such as a computer workstation, CRT display, and printer, represented by the workstation 234.

Information required to operate the network controller 214 will be stored in databases 226 and local memory 230 (e.g., either in RAM, ROM, or magnetic or optical Read/Write devices) at the cable headend 208 as well as in memory (RAM and/or ROM) within each subscriber's set top terminal 220. In the preferred embodiment, two-way communications between the network controller 214 and set top terminal 220 will occur over cable lines. Many other methods of communication, including those which do not require cables or wires, may be used with the present invention. Using two-way communication, interactive television programming can be accommodated through the network controller 214. In addition, the preferred network controller 214 will be able to access set top terminals 220 via phone lines for trouble shooting, special features or sophisticated reprogramming.

The network controller CPU 224 controls the interface, depicted at 211, between the network controller 214 and the signal processor 209. This interface 211 allows control information to flow or transfer between the two cable headend 208 components. Standard RS-232 or RS-422 links, an IEEE-488 bus or other interface media may be used. During standard operation, program control information is passed through this interface 211 to the network controller CPU 224 from the signal processor 209 (i.e., the program control information having been sent to the signal processor 209 over satellite from the operations center 202 with the RF program signals 205, not shown in Figure 5). The network controller CPU 224 processes the program control information based on data stored in the network control databases. This processing includes modifying the program control information to accommodate regional programming needs.

After processing, the network controller CPU 224 passes the program control information, including any modifications, back to the signal processor 209 for distribution over the cable system 200, via the cable distribution network 236. In this fashion, the network controller 214 provides programming and network control instructions to the set top terminals 220 through the signal processor 209.

The processing of program control information by the network controller CPU 224 can also make use of any data received by the network controller's control receiver 228. The control receiver 228 is a microprocessor-based device that receives "status reports" directly from the set top terminals 220. The status reports received by the control 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.

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

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.

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. PCT/US93/11708, entitled REPROGRAMMABLE TERMINAL FOR

SUGGESTING PROGRAMS OFFERED ON A TELEVISION PROGRAM DELIVERY SYSTEM, filed by the same assignee).

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.

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.

TABLE A

PM

	*Program name	*Program length	*Menu code	*Description	*Video
	Cheers	.5	E24	C	A
	Terminator	2.0	A33	Tx	S
	PrimeTime	1.0	D14	N	A
	Football Special	.5	B24	S	A

12:30 PM

	*Program name	*Program length	*Menu code	*Description	*Video
1	Simpsons	.5	E14 & C13	C	S
4	Football Game	3.0	B13	S	A
*					
*					
*					

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.

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 1 = YCTV™ 2 = Pay-Per-View 3 = Reg. TV	Unsigned Int
2	Event ID	Unsigned Int
3	Global Channel ID	Unsigned Int
4	Price (in Cents)	Unsigned Int
5	Start Time	HH:MM:SS
6	End Time	HH:MM:SS
7	Start Date	MM/DD/YY
8	End Date	MM/DD/YY
9	P-Icon	ASCIIZ
10	Name	ASCIIZ
11	Description	ASCIIZ

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

1 1234 2 50 03:00:00 15:00:00 08/25/93 08/27/93 pbs.pcx Sesame Street & Barney's Sesame Street and Barney Abstract
2 1234 3 50 20:00:00 22:00:00 08/25/93 08/25/93 t4.pcx Terminator 4 Terminator 4 Abstract

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 changed without physically swapping the ROM holding the 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 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.

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 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 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 250.

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 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. PCT/US93/11615, entitled DIGITAL CABLE HEADEND FOR CABLE TELEVISION DELIVERY SYSTEM, 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.

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. 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.

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 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

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 9a diagrams an out-of-band two-way data transmission system 280 wherein satellite feeds 282 are received at the cable headend 208 by a number of satellite receivers 284 and digital signal processing equipment 286. The 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 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. 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).

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.

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, may readily be accommodated by the embodiment shown in Figure 9a. Downstream services may include TV 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 306 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.

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 288. 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 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 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 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 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.

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 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 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 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 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. <u>Processing Information Received from Set top Terminals</u>

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.

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, 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 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 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 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 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 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 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 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 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.

8. <u>Overview of Software Routines</u>

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 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. 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 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, 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 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.

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 16 shows an example of a 30-day 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 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 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 example, during the 0800-1200 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. <u>Basic Advertisement Targeting Routine</u>

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 infomericals, with infomericals 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 426, 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 456.

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.

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 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 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) 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 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 Columns 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 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 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 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 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 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 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. PCT/US93/11606, entitled, ADVANCED SET TOP TERMINAL FOR CABLE TELEVISION DELIVERY SYSTEMS. 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. PCT/US93/11618, owned by the assignee of the present invention. 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.

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 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 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 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 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 menus 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 scope of the invention as defined in the following claims.

The embodiments of the invention in which an exclusive property or privilege is claimed are defined as follows:

1. A method for targeting advertising to at least one subscriber, the method comprising 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. The method of claim 1, wherein at lest one advertisement within an advertisement category is a promotion.

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

4. The method of claim 1, further comprising displaying the selected advertisement.

5. The method of claim 4, wherein the selected advertisement is displayed on a display of a television.

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

7. The method of claim 6, 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.

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

9. 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 comprising masking undesired advertisements out of the plurality of advertisements.

10. 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.

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

12. 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.

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

14. The method of claim 13, wherein the set top analyzes the programs watched data.

15. The method of claim 13, wherein a processor coupled to the set top receives the programs watched data and analyses the programs watched data.

16. The method of claim 15, further comprising:

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.

17. The method of claim 16, wherein processing the received status reports comprises:

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

reading at least one information filed 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.

18. The method of claim 17, further comprising updating the programs watched data based on the processed status reports.

19. The method of claim 15, wherein the set top transmits a set top status report that contains the programs watched data to the processor using a random access method.

20. The method of claim 19, wherein the random access method further comprises a CDMA/CD protocol.

21. The method of claim 15, wherein the set top transmits a set top status report that contains the programs watched data to the processor using a standard telephone line.

22. The method of claim 1, further comprising:

gathering demographic data related to the subscribers; and

correlating the demographic data with the categories of advertisements.

23. The method of claim 22, wherein gathering the demographic data comprises one of sending questionnaires to the subscribers, receiving demographic data via the set top and gathering demographic data from third parties.

24. The method of claim 22, wherein gathering demographic data comprises 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.

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

26. A system for targeting advertising to at least one subscriber comprising:

a means for gathering programs watched data from a subscriber;

a processor, operably connected to the gathering means, comprising 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.

27. The system of claim 26, wherein at least one advertisement within an advertisement category is an infomercial.

28. The system of claim 26, wherein at least one advertisement within an advertisement category is a promotion.

29. The system of claim 26, wherein the system is located at a cable headend.

30. The system of claim 26, wherein the system is located at an operations center.

31. The system of claim 26, 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.

32. The system of claim 26, wherein the processor further comprises:

a memory, operably connected to the receiver, wherein the memory stores video corresponding to the selected advertisement.

33. The system of claim 26, further comprising a display, wherein the transmitted advertisement is displayed.

34. The system of claim 33, wherein the display displays a menu of the selected advertisements from which the subscriber chooses an advertisement to watch.

35. The system of claim 33, further comprising:

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

36. The system of claim 33, wherein the processor is a network controller and the display is a set top terminal operably connected to a television.

37. The system of claim 36, wherein the transmitter transmits the selected advertisements on a television channel, and the set top terminal is capable of changing television channels.

38. The system of claim 36, 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.

39. The system of claim 36, 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.

40. The system of claim 26, 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.

41. The system of claim 40, wherein the programs watched matrix is arranged by program category and time slot and the analyzing means comprises:

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.

42. The system of claim 26, 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 comprising:

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.

43. The system of claim 26, wherein the means for gathering programs watched data is a set top terminal comprising memory wherein programs watched data is stored; and

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

44. The system of claim 43, 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 comprises:

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.

45. The system of claim 44, wherein the means for report processing comprises:

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.

46. The system of claim 45, further comprising a means for updating the programs watched data in the storing means with the polling response data.

47. The system of claim 26, further comprising:

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.

48. The system of claim 47, 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 comprises:

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.

49. A system for targeting advertising comprising:

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

a processor comprising:

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.

50. The system of claim 49, wherein the processor is a network controller.

51. A network controller for use in a cable television program delivery system for targeting advertising, the network controller comprising:

a means for gathering programs watched data from set top terminals;

a means for storing the gathered programs watched data;

a means for accessing the stored programs watched data;

a means for counting the accessed programs watched data to determine the frequency of programs watched by the set top terminals, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot;

a means for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each advertisement category includes at least one advertisement; and

a transmitter, wherein the transmitter transmits the set top terminal group information to the set top terminals in a control information stream that instructs the set top terminals in selecting advertisements for display during viewing of programs.

52. The network controller of claim 51, wherein the stored programs watched data is stored in a network control database.

53. A method of targeting advertisements comprising 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.

54. The method of claim 53, further comprising the step of storing the advertisement identified by the advertising company.

55. The method of claim 53, wherein at least one advertisement within an advertisement category is a promotion.

56. The method of claim 53, wherein at least one advertisement within an advertisement category is an infomercial.

57. The method of claim 53, further comprising 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.

58. The method of claim 57, wherein the switching is transparent to the viewer.

59. An apparatus for targeting advertising to a subscriber, the apparatus comprising:

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.

60. The apparatus of claim 59, wherein the information on advertisements identifies an infomercial.

61. The apparatus of claim 59, wherein the information on advertisements identifies a promotion.

62. The apparatus of claim 59, wherein the apparatus is a set top terminal.

63. The apparatus of claim 62, further comprising a receiver, wherein the receiver receives television signals.

64. The apparatus of claim 63, further comprising:

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

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

65. The apparatus of claim 63, wherein the receiver receives several channels of advertisements, and the apparatus further comprises:

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.

66. The apparatus of claim 59, further comprising:

means, operably connected to the correlating means, for storing the information on advertisements.

67. The apparatus of claim 66, wherein the storing means stores video corresponding to the selected advertisement.

68. The apparatus of claim 59, 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.

69. The apparatus of claim 68, wherein the correlating means generates an output further comprising:

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.

70. A method of targeting programs comprising 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.

71. A cable headend for targeting programs comprising:

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; and

means, operably connected to the selecting means, for storing programs, wherein the selected program is stored.

72. The apparatus of claim 71, further comprising:

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

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

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.

74. The apparatus of claim 73, wherein the video is television programming.

75. The apparatus of claim 73, wherein the video is advertisements.

76. A method for targeting advertising using subscriber information, comprising:

gathering programs watched data from a subscriber;

analyzing the programs watched data to determine a 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 a first advertisement and a second advertisement from the correlated advertising categories;

transmitting the first advertisement for display to a first subscriber and the second advertisement for display to a second subscriber; and

gathering advertisements watched data from the first and the second subscribers.

77. The method of claim 76, wherein the programs watched data comprises:

programs a viewer purchased and date/time of purchase;

channels the subscriber viewed and length of viewing; and

click stream data recorded during viewing by the subscriber.

78. The method of claim 76, wherein gathering advertisements watched data comprises gathering channel selected data during program breaks.

79. The method of claim 76, wherein analyzing the programs watched data further comprises:

determining how many subscribers watched a particular program;

determining a targeting group to assign the subscriber to;

determining peak viewing times for different program categories; and

determining advertising for advertisement spots.

80. The method of claim 76, wherein at least one advertisement within an advertisement category is a promotion.

81. The method of claim 76, wherein at least one advertisement within an advertisement category is an infomercial.

82. The method of claim 76, wherein analyzing programs watched data includes counting programs watched by the first and the second subscriber and further comprising arranging the counts in at least one programs watched matrix.

83. The method of claim 82, 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.

84. The method of claim 76, wherein the television terminal analyzes the programs watched data.

85. The method of claim 76, further comprising:

generating a polling request message that directs the television terminal to transmit a television terminal status report that contains the programs watched data and the advertisements watched data;

processing the received television terminal status reports to produce polling response data; and

storing the polling response data.

86. The method of claim 85, wherein processing the received television terminal status reports comprises:

demodulating the received status reports, producing demodulated television terminal data;

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

sorting each information field by television terminal identification number.

87. The method of claim 85, further comprising updating the programs watched data based on the processed television terminal status reports.

88. The method of claim 85, further comprising:

gathering demographic data related to the subscribers; and

correlating the demographic data with the categories of advertisements.

89. The method of claim 88, wherein gathering the demographic data comprises one or more of sending questionnaires to the subscriber, receiving demographic data via the television terminal and gathering demographic data from third parties.

90. The method of claim 89, wherein gathering demographic data from third parties comprise determining Area of Dominant Influence (ADI) and area code + 4 data.

91. The method of claim 89, wherein gathering demographic data comprises 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.

92. The method of claim 89, wherein the television terminal provides the television terminal status report using a random access method and further including using a collision detection/collision avoidance protocol.

93. The method of claim 85, wherein the television terminal transmits a television terminal status report that contains the programs watched data and the advertisements watched data to a remote site using a standard telephone line.

94. The method of claim 93, wherein the remote site is a cable television headend.

95. The method of claim 93, wherein the remote site is a broadcast television operations center.

96. The method of claim 76, wherein the method for targeted advertising is performed at one of a national and a regional operations center.

97. The method of claim 76, wherein the method for targeted advertising is performed at a cable headend site.

98. The method of claim 76, wherein the method for targeted advertising is performed at a remote site that functions as an operations center and a cable headend site.

99. The method of claim 76, wherein the subscriber is assigned to a group based on the analyzed programs watched data and the advertisements watched data.

100. The method of claim 99, further comprising:

providing a plurality of advertisements at a television terminal used by the subscriber for viewing programs, the plurality of advertisements provided over a program channel and one or more feeder channels; and

generating a switching plan, wherein the switching plan instructs the television terminal to select one of the program channel and the at least one feeder channel during a program break occurring in a broadcast of a program on the program channel.

101. The method of claim 100, wherein the selection of the program channel and the feeder channel is based on the group assignment.

102. A method for targeting advertisements to television viewers based on television viewer information, comprising:

receiving gathered television viewer information;

analyzing the received viewer information;

correlating the analyzed viewer information to advertisement categories;

selecting a first advertisement from a first advertisement category;

selecting one or more subsequent advertisements from one of the first advertisement category and one or more subsequent advertisement categories;

providing the selected first advertisement and the one or more subsequent advertisements to television terminals;

providing a switching plan that directs the television terminals to display one of the provided first advertisement and the one or more subsequent advertisements during an advertisement spot in a television program.

103. The method of claim 102, wherein the providing step comprises providing the selected advertisement and the one or more subsequent advertisements for storage at the television terminals.

104. The method of claim 102, wherein the providing step comprising:

providing the selected first advertisement on a program channel displaying the television program; and

providing the selected one or more subsequent advertisements on one or more feeder channels.

105. The method of claim 102, wherein the television viewer information comprises television programs watched information.

106. The method of claim 102, wherein the television viewer information comprises advertisements watched information.

107. The method of claim 102, wherein the television viewer information comprises viewer demographic information.

108. The method of claim 102, wherein the television viewer information is gathered by the television terminals.

109. The method of claim 102, wherein the television viewer information is gathered by third parties.

110. The method of claim 109, wherein the third parties include cable television headends.

111. The method of claim 102, wherein the switching plan is based on the analyzed viewer information.

112. The method of claim 102, wherein generating the switching plan comprises generating a national switching plan.

113. The method of claim 112, further comprising:

generating a local switching plan, whereby local advertisements are selected for display at the television terminals;

modifying the national switching plan based on the local switching plan; and

providing the modified national switching plan to selected ones of the television terminals.

114. A method for targeting advertisements to television viewers in a television program delivery system, comprising:

determining advertisement categories;

mapping television terminals to the determined advertisement categories, wherein one or more of the television terminals are assigned to groups for one or more of the advertisement categories; and

mapping television programs to the groups.

115. The method of claim 114, further comprising:

ranking two or more advertisements within one or more of the advertisement categories; and

selecting for display, advertisements that have a higher ranking.

116. The method of claim 115, further comprising:

determining a number of available channels for display of advertisements;

assigning the advertisements to the available channels in order of the advertisement's ranking until all available channels are assigned an advertisement.

117. A network controller for remotely managing account and billing information over a cable television distribution network characterized by:

a means for creating a polling request message;

a means for transmitting the polling request message to set top terminals, wherein the polling request message directs each set top terminal to initiate upstream data transmission of set top terminal status reports over the cable distribution network;

a means for receiving the set top terminal status reports, wherein the set top terminal status reports include programs watched data;

a means for storing network control data, wherein the stored network control data includes programs watched data and price category data;

a means for correlating the programs watched data with the price category data to produce price correlations; and

a means for generating at least one billing report based on the produced price correlations; and

a means for transmitting the billing reports to the set top terminals for display on a television screen.

118. A method for remotely managing account and billing information over a cable television distribution network for a plurality of set top terminals, the method characterized by the steps of:

creating a polling request message;

transmitting the polling request message to the set top terminals, wherein the polling request message directs each set top terminal to initiate upstream data transmission of set top terminal status reports over the cable distribution network;

receiving the set top terminal status reports, wherein the set top terminal status reports include programs watched data;

storing network control data, wherein the stored network control data includes programs watched data and more than one price category;

correlating the programs watched data with the price categories, wherein price correlations are produced; and

generating at least one billing report based on the produced price correlations; and

transmitting the billing reports to the set top terminals for display on a television screen.

119. A network controller for use with a cable headend capable of remotely targeting categories of advertisements to groups of set top terminals in a cable television system, the network controller characterized by:

a means for gathering programs watched data for each set top terminal in the cable television system;

a means for counting the gathered programs watched data to determine the frequency of programs watched by each set top terminal in the cable television program delivery system, wherein the programs watched counts are compiled by program category and time slot;

a means for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each category advertisements includes one or more advertisements available for targeting to the set top terminals; and a means for transmitting the set top terminal group information to the set top terminal in a control information stream that instructs the set top terminal in selecting targeted advertisements for display during viewing of programs.

120. A network controller for remotely monitoring a plurality of set top terminals in a cable television system, the network controller characterized by:

a means for receiving a program control information signal from digital signal processing equipment, wherein the digital signal processing equipment is co-located with the receiving means at the cable headend and the program control information signal is received by the digital signal processing equipment from a remotely located source;

a means for storing network control data, wherein the stored network control data includes data on television programs;

a means for accessing the stored network control data;

a means for generating a control information stream using the received program control information signal and the accessed network control data, wherein the received program information signal carries data on packaged programs or menu content and whereby information fields of the control information stream are created by modifying the data on packaged programs or menu content based on the stored network control data so that the created information fields contain data on modified packages of programs or modified menu content, including program category and menu assignment information;

a means for transferring the control information stream to the digital signal processing equipment, whereby the control information stream is subsequently distributed by the digital signal processing equipment to multiple set top terminals;

a means for gathering programs watched data for each set top terminal in the cable television program delivery system, wherein the gathered programs watched data is stored in the storing means so that it can be accessed by the accessing means;

a means for counting the accessed programs watched data to determine the frequency of programs watched by each set top terminal in the cable television program delivery system, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot; and

a means, connected to the generating means, for creating set top terminal group information indicating a group assignment for each set top terminal by correlating the programs watched counts with the categories of advertisements, wherein each advertisements category includes advertisements available for targeting to the set top terminals, and wherein the created set top terminal group information is included as data in the control information stream.

121. The network controller of claim 120, characterized in that the gathering means includes:

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

a means for receiving the set top terminal status reports, wherein the received set top terminal status reports contain programs watched information;

a means for processing the received set top terminal status reports to produce polling response data; and

a means for temporarily storing the polling response data.

122. The network controller of claim 121, characterized in that the processing means 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 for accumulating the sorted information fields for each set top terminal, wherein the accumulated sorted information fields produce the polling response data.

123. The network controller of claim 121 further characterized by a means for updating the programs watched data in the storing means with the polling response data.

124. The network controller of claim 121, characterized in that the counting means includes:

a means for reading the programs watched counts, wherein a separate programs watched counts is assigned to different time slots in a day for each program category;

a means for sorting the programs watched counts from highest to lowest;

a means for matching the sorted programs watched counts with the advertisements categories, wherein counts of the advertisements categories are thereby produced.

125. A network controller for remotely monitoring a plurality of set top terminals in a cable television system, the network controller characterized by:

a means for receiving a program control information signal from digital signal processing equipment, wherein the digital signal processing equipment is co-located with the receiving means at the cable headend and the program control information signal is received by the digital signal processing equipment from a remotely located source;

a means for storing network control data, wherein the stored network control data includes data on television programs;

a means for accessing the stored network control data;

a means for generating a control information stream using the received program control information signal and the accessed network control data, wherein the received program information signal carries data on packaged programs or menu content and whereby information fields of the control information stream are created by modifying the data on packaged programs or menu content based on the stored network control data so that the created information fields contain data on modified packages of programs or modified menu content, including program category and menu assignment information;

a means for transferring the control information stream to the digital signal processing equipment, whereby the control information stream is subsequently distributed by the digital signal processing equipment to multiple set top terminals;

a means for gathering programs watched information from groups of the set top terminals;

a means for storing the gathered programs watched information and the individual advertisements arranged by advertisement category;

a means for sorting the stored programs watched information by a plurality of program categories;

a means for ranking the plurality of program categories by frequency of programs watched in each category, wherein program categories with more programs watched are ranked higher than program categories with less programs watched;

a means for ranking the stored advertisement categories, wherein the ranked program categories are correlated with the stored advertisement categories;

a means for selecting a set of highest ranked advertisement categories;

a means for choosing individual advertisements from the set of highest ranked advertisement categories; and

a means, connected to the generating means, for packaging the chosen individual advertisements for transmission to the set top terminals for display, wherein packages of targeted advertisements are produced and included in the control information stream and sent to the transferring means.

126. The network controller of claim 125, characterized in that the gathering means includes:

a means for receiving polling responses from the set top terminals, wherein the polling responses are transmitted upstream over the cable distribution network in response to a polling request message that is sent to the set top terminals in the control information stream generated by the generating means; and

a means for processing the received polling responses, wherein data on the programs watched information for the set top terminals is extracted.

127. The network controller of claim 126, characterized in that the gathering means further includes a means for updating previously stored programs watched information in the storing means using the extracted data.

128. A network controller for remotely monitoring a plurality of set top terminals in a cable television system, the network controller characterized by:

a means for receiving a program control information signal from digital signal processing equipment, wherein the digital signal processing equipment is co-located with the receiving means at the cable headend and the program control information signal is received by the digital signal processing equipment from a remotely located source;

a means for storing network control data, wherein the stored network control data includes data on television programs;

a means for accessing the stored network control data;

a means for generating a control information stream using the received program control information signal and the accessed network control data, wherein the received program information signal carries data on packaged programs or menu content and whereby information fields of the control information stream are created by modifying the data on packaged programs or menu content based on the stored network control data so that the created information fields contain data on modified packages of programs or modified menu content, including program category and menu assignment information;

a means for transferring the control information stream to the digital signal processing equipment, whereby the control information stream is subsequently distributed by the digital signal processing equipment to multiple set top terminals;

a means for receiving set top terminal status reports, wherein the set top terminal status reports are transmitted by the set top terminals upstream over the cable distribution network in response to a polling request message that is sent to the set top terminals in the control information stream generated by the generating means and wherein the received set top terminal status reports include programs watched data that is stored in the storing means for use by the generating means in generating the control information stream;

a means, connected to the accessing means, for correlating the programs watched data with price category data stored in the storing means to produce price correlations; and

a means for generating at least one billing report based on the produced price correlations, wherein the billing reports are included in the control information stream and sent to the transferring means for subsequent distribution to the set top terminals.

129. An apparatus for targeting advertising to at least one subscriber, comprising by:

means for gathering programs watched data from a subscriber; and

a processor operably connected to the gathering means and including:

means for analyzing the gathered programs watched data by determining the frequency of programs watched by the subscriber or the viewing habits of the subscriber;

means for correlating the analyzed programs watched data with at least one advertisement by comparing the programs watched data with advertisements;

means for selecting at least one correlated advertisement.

130. The apparatus of claim 129 further including a display for displaying the selected advertisement.

131. The apparatus of claim 129 further including a display for displaying a menu listing the selected advertisement.

132. The apparatus of claim 129, 130 or 131, further including a receiver operably connected to the processor for receiving television signals.

133. The apparatus of any one of claims 129 to 132, wherein the gathering means and processor are located in a set top terminal.

134. The apparatus of claim 133, wherein the set top terminal includes a memory for storing video corresponding to the selected advertisement.

135. The apparatus of claim 129 further including a receiver, wherein the receiver receives several channels of advertisements.

136. The apparatus of claim 135 further including:

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

137. The apparatus of claim 135, wherein the receiver also receives directions which instruct a set top terminal to video corresponding to the selected advertisement, and further including a memory for storing the video corresponding to the selected advertisement.

138. The apparatus of claim 129, wherein the gathering means and processor are located in a cable headend, a network controller or an operations center.

139. The apparatus of claim 138 further including a transmitter for transmitting the selected advertisement.

140. The apparatus of claim 129, wherein the analyzing means creates a viewer file.

141. The apparatus of claim 140, wherein the frequency of programs watched by the subscriber is arranged in at least one programs watched matrix.

142. The apparatus of claim 141, wherein the programs watched matrix is arranged by program category and time slot and the analyzing means includes:

means for reading the programs watched matrix; and

means for sorting programs watched counts within a time slot from highest to lowest.

143. The apparatus of any one of claims 129 to 142, wherein the correlated advertisement is a category of advertisements and the category of advertisements includes at least one advertisement.

144. The apparatus of claim 143, wherein the category of advertisements includes at least one promotion or at least one infomercial.

145. The apparatus of any one of claims 129 to 144, wherein the correlated advertisement is an infomercial or a promotion.

146. The apparatus of any one of claims 129 to 144, wherein the gathering means gathers programs watched data from a plurality of subscribers and the processor processes the programs watched data gathered from the plurality of subscribers.

147. The apparatus of any one of claims 129 to 146 further including:

means for developing a program line-up based on the analyzed programs watched data, wherein the program line-up includes advertisements; and means operably connected to the developing means for transmitting the program line-up to the subscriber.

148. The apparatus of claim 147 further including means for gathering marketing data, wherein the analyzing means also analyses the marketing data.

149. The apparatus of any one of claims 129 to 148 further including:

a memory operably connected to the processor for storing at least one program for selection; and

means operably connected to the memory for choosing at least one stored program based on the analyzed data.

150. The apparatus of claim 149 further including means operably connected to the choosing means for displaying the chosen program to the subscriber.

151. The apparatus of claim 129, wherein the processor includes a first and a second processor, the first processor including the analyzing means, and the second processor including the correlating means and the selecting means.

152. The apparatus of claim 129, wherein the processor includes a first and second processor, the first processor including the analyzing means and the correlating means, and the second processor including the selecting means.

153. The apparatus of claim 151 or 152, wherein the first processor is in a network controller or in an operation center and the second processor is in a set top terminal.

154. The apparatus of claim 129, wherein the processor polls at least one set top terminal to monitor programs being watched by the subscriber for the occurrence of commercial breaks, further including:

means for polling at least one set top terminal; wherein the transmitter transmits at least one selected advertisement in response to the occurrence of commercial breaks determined by the polling means.

155. The apparatus of claim 129, wherein the gathering means is a set top terminal comprising a memory and the memory stores programs watched data, and the processor further includes means for accessing the stored programs watched data from the set top terminal.

156. The apparatus of claim 155, wherein the processor polls the set top terminal for status reports, the accessing means accesses the status reports sent by the set top terminal, and the accessing means further includes:

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

means for status report processing, wherein the status reports are processed to produce polling response data; and

means for temporarily storing the polling response data.

157. The apparatus of claim 156, wherein the means for report processing includes:

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

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;

means for sorting each information field in the status reports by a set top terminal identification number; and temporary memory means, wherein each sorted information field is accumulated and each accumulated sorted information field produces the polling response data.

158. The apparatus of claim 156 further including:

means for storing the programs watched data; and

means for updating the programs watched data in the storing means with the polling response data.

159. The apparatus of claim 129 further including means for gathering demographic data from subscribers, wherein the means for correlating also uses the gathered demographic data.

160. The apparatus of claim 159, wherein a first group of subscribers forms a statistically significant number of subscribers, the means for gathering demographic data gathers demographic data from the first group of subscribers, the means for gathering programs watched data gathers programs watched data from the first group of subscribers and a second group of subscribers, and the apparatus further includes means for generating a simulated demographic profile of the second group of subscribers by comparing the programs watched data of the second group of subscribers with the gathered demographic data and the programs watched data of the first group of subscribers.

161. The apparatus of claim 129, wherein the processor further includes:

means for counting the gathered programs watched data to determine the frequency of programs watched by the subscribers, wherein the programs watched counts are arranged in at least one programs watched matrix by program category and time slot; and means for creating subscriber group information indicating a group assignment for each subscriber by correlating the programs watched counts with at least one advertisement; and

wherein the transmitter transmits the subscriber group information in a control information stream.

162. The apparatus of claim 161 further including a set top terminal, wherein the control information stream instructs the set top terminal in selecting advertisements for display during viewing of programs.

163. The apparatus of any one of claims 129 to 162, wherein at least one advertisement is a promotion or an infomercial.

164. The apparatus of any one of claims 129 to 163, wherein at least one advertisement includes information on advertisements but not video corresponding to the information on advertisement.

165. The apparatus of any one of claims 129 to 164, wherein at least one advertisement includes information on advertisements and video corresponding to the information on advertisements.

166. The apparatus of claim 165, wherein the information on advertisements identifies an infomercial or a promotion.

167. The apparatus of claim 166, wherein the information on advertisements identifies the individual advertisements by name or type of advertisement.

168. The apparatus of claim 165, wherein the information on advertisements includes directions for inserting advertisements into a program line-up.

169. A method of targeting advertisements, comprising:

gathering programs watched data from a subscriber;

analyzing gathered programs watched data by determining the frequency of programs watched by the subscriber or the viewing habits of the subscriber;

correlating the analysed programs watched data with at least one advertisement by comparing the programs watched data with advertisements; and

selecting an advertisement based on programs watched data.

170. The method of claim 169 further including the step of transmitting the selected advertisement for display to the subscriber.

171. The method of claim 169 or 170 further including the step of storing the selected advertisement.

172. The method of any one of claims 169 to 171, wherein the gathered programs watched data is analysed to determine a viewer profile for the subscriber.

173. The method of any one of claims 169 to 172, wherein the advertisement is a promotion or an infomercial.

174. The method of any one of claims 169 to 173, wherein the advertisement is information on advertisements, but no video corresponding to the information on advertisements.

175. The method of any one of claims 169 to 173, wherein the advertisement is information on advertisements and video corresponding to the information on advertisements.

176. The method of any one of claims 169 to 173, wherein the advertisement is a category of advertisements, and the category of advertisements includes at least one advertisement.

177. The method of claim 176, wherein the category of advertisements includes at least one promotion or at least one infomercial.

178. The method of any one of claims 169 to 177 further including the steps of:

receiving several channels of advertisements;

wherein the selected advertisement is within a received channel; and switching to the channel containing the selected advertisement.

179. The method of claim 178, wherein the step of switching is transparent to the viewer.

180. The method of any one of claims 169 to 179 further including the steps of:

storing at least one program for selection at a remote location; choosing at least one stored program based on the analyzed data; and transmitting at least one chosen program from the remote location to the subscriber.

181. An apparatus for providing digital program signals to subscriber locations, capable of inserting local availability signals and selecting digital program signals received from outside sources, comprising:

a means for receiving digital program signals;

digital logic circuitry, connected to the receiving means, wherein digital program signals may be inserted and wherein digital programs may be selected;

a processor, operably connected to the digital logic circuitry, wherein insertion of local availability signals are controlled and wherein the selection of digital program signals are controlled; and

means for sending programs, operably connected to the digital logic circuitry, wherein digital program signals are sent to subscriber locations, and wherein digital program signals that have been inserted or selected are sent to subscriber locations.

182. The apparatus of claim 181 wherein the digital logic circuitry comprises a combiner, wherein digital program signals are combined.

183. The apparatus of claim 181 wherein the digital logic circuitry comprises a local inserter, wherein local availability signals are inserted to be sent to subscriber locations.

184. The apparatus of claim 181 wherein the sending means comprises a serializer wherein the serializer serializes digital program signals.

185. The apparatus of claim 181 wherein the receiving means comprises an integrated receiver decoder.

186. An apparatus for use in a program delivery system having a plurality of program signals, wherein the apparatus receives directions for local insertion from a remote source, comprising:

a receiver, wherein directions for local insertion from a remote source are received;

a database, containing information on a plurality of programs;

a processor, operably connected to the receiver and database, wherein the directions from the remote source are processed to generate a processed signal;

a local inserter, connected to the processor, wherein local programming is inserted into a digital program signal using the processed signal from the processor; and

a modulator, wherein a plurality of program signals including the program signal carrying the inserted local programming is modulated.

187. The apparatus of claim 206 further comprising a signal processor, operably connected to the modulator, wherein the signal processor receives a plurality of program signals from the remote source.

188. The apparatus of claim 206 wherein the receiver receives selecting directions on a subset of programs from the plurality of programs to be selected from a remote source, the processor processes the selecting directions, and wherein the local inserter comprises:

digital logic, wherein the digital logic selects programs using the processed directions, and wherein the modulator modulates the selected programs.

189. The apparatus of claim 206 wherein the local programming is advertising and the directions for local insertion are directions for inserting advertisements.

190. An apparatus for a program signal delivery system wherein digital program signals are used comprising:

a receiver, wherein program signals are received;

a demultiplexor, connected to the receiver, wherein program signals are demultiplexed for processing;

digital logic, connected to the demultiplexer, wherein demultiplexed signals are processed;

a control processing unit, operably connected to the digital logic, wherein the processing of the demultiplexed signals in the digital logic is controlled;

multiplexor, connected to the digital logic, wherein processed signals are multiplexed; and

a modulator, connected to the multiplexer, wherein multiplexed signals are modulated and wherein the modulated signal is forwarded to subscribers.

191. The apparatus of claim 190 wherein the processing of signals includes local insertion of program signals.

192. The apparatus of claim 190 wherein the processing of signals includes selection of certain demultiplexed signals.

193. An apparatus for use with a network having subscribers that uses two-way communications allowing upstream data transmission from the subscriber comprising:

a two-way communication system comprising

one or more diplex filters, wherein the diplex filters facilitate upstream data transmissions;

a signal processor, operably connected to the two-way communication system, wherein signals are processed before being sent downstream to subscribers;

an RF combiner, connected to the two-way communication system, wherein upstream data transmissions are received from the two-way communication system;

a network controller, connected to the RF combiner, wherein upstream data transmissions are received from the RF combiner; and

wherein program control information is used, and the network controller comprises a processing unit, wherein program control information is processed.

194. The apparatus of claim 193 wherein the network controller comprises memory and a processing unit, wherein programs watched information received from the RF combiner is processed by the processing unit and stored in the memory.

195. The apparatus of claim 193 herein the network controller comprises:

a memory, wherein an interactive game program is stored; and

a processor, operably connected to the memory, wherein the processor executes an interactive game program.

196. A computer system for a network which controls video and audio delivery to subscriber locations, comprising:

a network connecting subscriber locations;

a plurality of terminals, connected to the network, to allow subscribers at subscriber locations to access the network;

a processor, connected to the network and remotely located from the subscriber locations, wherein the processor generates control signals to control the distribution of video and audio signals to the plurality of terminals located at subscriber locations; and

signal processing equipment, which communicates with the processor through a non-network connection, adapted to output audio and video signals to a plurality of terminals.

197. The computer system of claim 196 wherein subscribers interact with other subscribers on the network and the processor processes subscriber requests to interact with other subscribers.

198. The computer system of claim 197 wherein the means to allow subscribers to interact with other subscribers comprises an interactive game, wherein more than one subscriber plays the same game.

199. The computer system of claim 196 further comprising a receiver, operably connected to the processor, wherein the digital video and audio signals are television programs received via satellite transmission by the receiver.

200. A network controller for use in a cable system for choosing text overlays, wherein a subscriber may request a text overlay representing the audio associated with the video to a program, wherein a set top terminal transmits the subscriber requests comprising:

a means for receiving subscriber requests;

a processor to identify the subscriber request for text overlays from a plurality of subscriber requests;

a means for choosing the corresponding text overlay from a plurality of text overlays; and

a means for transmitting text overlay to the set top terminal for display.

201. The network controller of claim 84 further comprising:

means for receiving control signals with text overlay;

means for processing the received control signals for text overlay.

202. The network controller of claim 200 wherein programs may have sign language video associated with the program wherein the sign language video is displayed.

203. The network controller of claim 202 wherein the sign language video is transmitted on a separate channel to the subscriber.

204. A system for choosing audio channels, wherein more than one audio channel may exist for a video, comprising:

a signal processor, wherein a plurality of digital video signals and audio channels are received, and wherein a received video signal may have multiple audio channels;

a controller, wherein the controller selects one or more received audio channels, wherein one or more audio channels that correspond to a received video are selected;

a distribution system, operably connected to the controller, wherein the selected audio channels are distributed and a listener receives one or more audio channels; and

wherein a program control information signal is used by the controller to select one or more received audio channels.

205. The system of claim 204 wherein the program control information signal is received by the signal processor, further comprising an interface connecting the signal processor and the controller, wherein the program control information signal is passed between the signal processor and controller using the interface.

206. The system of claim 204 wherein each audio channel received by the signal processor for a particular video is in a different language and the controller selects audio channels based on language desired.

207. An apparatus that gathers information related to television viewing habits, comprising:

a first processor that provides information related to which television programming was viewed;

a first memory coupled to the first processor that stores the information;

a second processor that receives the information from the first memory and arranges the information as programs watched data; and a second memory coupled to the second processor that stores the programs watched data.

208. The apparatus of claim 207, wherein the information includes date and time information, and program information and wherein the programs watched data is stored as a programs watched matrix.

209. The apparatus of claim 208, wherein the program information includes one or more of program category, program title information, and program duration.

210. The apparatus of claim 209, wherein the program category includes one or more of sports, movies, children's shows, news shows and pay-for-view events, and wherein the date and time information is the date and time the program was viewed.

211. The apparatus of claim 209 or 210, wherein the programs watched matrix includes columns and rows, the columns listing the time and date information and the rows listing the program information.

212. The apparatus of claim 211, wherein the time and date information is listed in blocks of four hours.

213. The apparatus of any one of claims 207 to 212, wherein the information is provided from a plurality of set top terminals that receive programming from a television delivery system, and wherein each of the plurality of set top terminals is identified by a unique identification number.

214. The apparatus of claim 213, wherein each particular set top terminal has an associated programs watched matrix, and programs

watched data for a particular set top terminal is arranged in the associated programs watched matrix.

215. The apparatus of claim 214, wherein as additional television programming is viewed, the programs watched matrix for the set top terminal is updated.

216. The apparatus of any one of claims 207 to 215, wherein the first processor and the first memory are contained in a set top terminal.

217. The apparatus of any one of claims 207 to 216, wherein the second processor and the second memory are contained in a cable television headend.

218. The apparatus of any one of claims 207 to 217, wherein the second processor and the second memory are contained in an operations center.

219. The apparatus of any one of claims 207 to 218, wherein the second processor and the second memory are contained in a set top terminal.

220. The apparatus of any one of claims 207 to 219, wherein the second processor obtains the information by sending a polling message to a set top terminal connected to the second processor, the set top terminal, in response to the polling message, sending the information to the second processor.

221. An apparatus that gathers programs watched data, comprising:

a plurality of terminals connected to corresponding televisions and to a television program delivery system, each of the terminals including a memory that stores program access information; and a receiver coupled to the plurality of terminals, the receiver receiving the program access information, wherein the program access information is stored as programs watched data.

222. The apparatus of claim 221, wherein the programs watched data is stored in a programs watched matrix, each of the terminals being assigned a unique programs watched matrix.

223. The apparatus of claim 222, wherein each of the terminals is assigned a unique address used to identify the terminal and the associated programs watched matrix.

224. The apparatus of claim 222 or 223, wherein each of the terminals is assigned a group identification, the group identification common to at least two terminals.

225. The apparatus of claim 222, 223 or 224, wherein the programs watched matrix includes time of day a program is watched and a program category for the program.

226. The apparatus of claim 225, wherein the program categories include one or more of children's programs, news programs, sports programs, pay-for-view, and movies.

227. The apparatus of any one of claims 222 to 226, wherein the programs watched matrix includes a time of day a program is watched and a program title.

228. The apparatus of any one of claims 222 to 227, wherein the unique programs watched matrix is updated as additional program access information is provided by the associated terminal.

229. The apparatus of claim 228, wherein the programs watched data is recorded as counts, the counts corresponding to the number of times a program category is watched at the associated terminal.

230. The apparatus of any one of claims 221 to 229, further comprising a plurality of databases, wherein each of the plurality of databases receives information from programs watched matrices, the information including the terminal address, group identifier, and program counts, the databases including one or more of viewer profile, account billing, program scheduling and advertisement scheduling databases.

231. The apparatus of any one of claims 221 to 230, further comprising a controller in the television program delivery system, the controller coupled to the plurality of terminals, the controller issuing a message directing each of the terminals to provide the program access information, wherein the message is a polling request message, the polling request message sent over one of a cable television cable and a telephone line and a response message is returned over one of the cable television cable and the telephone line.

232. The apparatus of claim 231, wherein the polling request message is a cyclic polling message, and wherein the cyclic polling message is one of roll-call polling, hub polling and token-passing polling.

233. The apparatus of any one of claims 231 or 232, wherein the program access information is provided in the response message, the response message including:

a leading flag;

an address field including the address of the terminal;

a subscriber region designation that includes a geographical region in which the terminal is located;

a terminal identifier that uniquely identifies the terminal;

an information field that includes a command to provide the program access information; and

a trailing flag.

234. The apparatus of claim 231, 232 or 233, wherein the plurality of terminals provide the program access information using a random access method.

235. The apparatus of any one of claims 231 to 234, wherein the controller is located in a cable headend of the television program delivery system.

236. The apparatus of any one of claims 231 to 235, wherein the controller is located in an operations center of the television program delivery system.

237. The apparatus of any one of claims 221 to 236, wherein individual terminals are grouped in one of a plurality of terminal groups, and wherein the programs watched data is stored as group programs watched data.

238. A system that gathers programs watched data in a broadcast television delivery system, comprising:

means for gathering programs watched data from one or more set top terminals in a broadcast television delivery system;

a databases that stores the gathered programs watched data;

means for accessing the stored programs watched data;

means for counting the accessed programs watched data to determine programs watched counts corresponding to the frequency of programs watched by the one or more set top terminals in the broadcast television delivery system, wherein the programs watched counts are arranged in at least one programs watched matrix;

means for creating set top terminal group information indicating a group assignment for a set top terminal by correlating the programs watched counts with categories of videos, wherein the video categories include videos available for sending to the set top terminal; and

means for transmitting the set top terminal group information to the set top terminal in a control information stream that instructs the set top terminal in selecting the videos for display.

239. The system of claim 238, further comprising:

means for extracting data from an information field of a program control information signal, wherein the extracted data includes program information;

means for creating a polling request message that directs a set top terminal to initiate transmission of a set top terminal status reports;

means for processing the received set top terminal status reports to produce polling response data; and

means for integrating the polling response data with the extracted data from the information field of the program control information signal.

240. The system of claim 239, wherein the creating means comprises:

means for reading database files using the access means, wherein the database files are relationally keyed to one another through a set of set top terminal identification numbers individually unique to a particular set top terminal;

means for formatting the polling request message, wherein the formatted polling request message includes at least one set top terminal identification number; and

means for enabling at least one polling command bit in the formatted polling request message, wherein the enabled polling command bit commands transmission of the set top terminal status reports.

241. The system of claim 239 or 240, wherein the processing means comprises:

control receiving means for demodulating the received set top terminal status reports;

means for reading at least one information field in the received set top terminal status reports, wherein the information field includes the programs watched data;

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

temporary memory means for accumulating the sorted information fields for the set top terminals, wherein the accumulated sorted information fields produce the polling response data.

242. The system of claim 241, further comprising connection means for linking the processing means to the database, wherein the polling response data may be stored in the database, updating the programs watched matrix.

243. The system of claim 241 or 242, wherein the control receiving means comprises a Radio Frequency demodulator for receiving upstream data transmissions from the set top terminals.

244. The system of claim 241, 242 or 243, wherein the control receiving means comprises a telephone for receiving data transmissions from the set top terminals over telephone lines.

245. An apparatus for use in a cable television program delivery system, the cable television delivery system coupled to one or more set top terminals, the one or more set top terminals capable of generating programs watched data, the apparatus comprising:

means for gathering programs watched data for a set top terminal;

means for storing the gathered programs watched data in at least one database;

means for accessing the stored programs watched data; and

means for counting the accessed programs watched data to determine the frequency of programs watched by the top terminal, wherein the programs watched counts are arranged by program category and time.

246. The apparatus of claim 245, further comprising:

means for creating set top terminal group information indicating a group assignment for the set top terminals; and

means for transmitting the set top terminal group information to the set top terminals in a control information stream capable of polling the set top terminals, and wherein the gathering means comprises:

means for generating a polling request message that directs the set top terminals to initiate transmission of a set top terminal status report;

means for receiving the set top terminal status report, wherein the received set top terminal status report contains programs watched information;

means for processing the received set top terminal status reports to produce polling response data; and

means for storing the polling response data.

247. The apparatus of claim 246, wherein the processing means comprises:

means for demodulating the received set top terminal status report;

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

means for sorting the at least one information field in received set top terminal status reports by set top terminal identification number; and

memory means for accumulating the sorted information fields for the at least one set top terminal, wherein the accumulated sorted information fields produce the polling response data.

248. The apparatus of claim 246 or 247, further comprising means for updating the programs watched data in the storing means with the polling response data, and wherein the counting means comprises:

means for reading the programs watched counts, wherein a separate programs watched count is assigned to different time slots in a day for a program category; and

means for sorting the programs watched counts from highest to lowest.

249. An apparatus that gathers programs watched information from set top terminals in a broadcast television program delivery system, comprising:

means for gathering the programs watched information;

means for storing the gathered programs watched information;

means for sorting the stored programs watched information by a plurality of program categories; and

means for ranking the plurality of program categories by frequency of programs watched in each category, wherein program categories with more programs watched are ranked higher than program categories with fewer programs watched.

250. The apparatus of claim 249, wherein the set top terminals have the capability of receiving polling request messages and transmitting polling responses, the polling responses including the programs watched information, and wherein the gathering means comprises:

means for generating the polling request messages, wherein the polling request messages are transmitted to the set top terminals, requesting the set top terminals to return the polling responses;

a receiver that receives the polling responses from the set top terminals; and

means for processing the received polling responses, wherein data on the programs watched information for the set top terminals is extracted and stored.

251. The apparatus of claim 249 or 250, wherein the gathered programs watched information has been previously stored in the storing means, and wherein the gathering means comprises means for updating the programs watched information in the storing means using the extracted data.

252. The apparatus of claim 249, 250 or 251, wherein the set top terminals transmit the programs watched information using CSMA/CD protocols.

253. The apparatus of any one of claims 249 to 252, wherein the programs watched information is transmitted over a cable television cable.

254. The apparatus of any one of claims 249 to 253, wherein the programs watched information is transmitted over a telephone line.

255. A method for gathering programs watched information, comprising:

gathering programs watched data for set top terminals in a broadcast television program delivery system;

storing the gathered programs watched data in a database;

accessing the stored programs watched data; and counting the accessed programs watched data to determine the frequency of programs watched by the set top terminals, wherein programs watched counts are arranged in a programs watched matrix.

256. The method of claim 255, wherein the set top terminals transmit set top terminal status reports in response to a polling request message produced using a program control information signal, further comprising:

creating the polling request message;

transmitting the polling request message to the set top terminals, wherein the polling request message directs set top terminals to initiate upstream data transmission of the set top terminal status reports over a cable distribution network; and

receiving the set top terminal status reports, wherein the set top terminal status reports include the programs watched data.

257. A processor in a cable television program delivery system, the delivery system providing programming to subscriber terminals, comprising:

means for gathering programs watched data for the terminals;

a database that stores the gathered programs watched data;

means for accessing the stored programs watched data; and

means for counting the accessed programs watched data to determine the frequency of programs watched at the terminals, wherein programs watched counts are arranged in a programs watched matrix.

258. The processor of claim 257, further comprising:

means for creating terminal group information indicating group assignments for the terminals by correlating the programs watched counts with categories of data, wherein the data categories include demographic data, and wherein the data categories are available for providing programming to the terminals; and means for transmitting the terminal group information to the terminals in a control information stream that instructs the terminals in selecting the programming, wherein the processor polls the terminals to direct the terminals to provide the programs watched data.

259. The processor of claim 158, wherein the gathering means comprises:

means for generating a polling request message that directs the terminals to initiate transmission of terminal status reports;

a receiver that receives the terminal status reports, wherein the received set terminal status reports contain programs watched information;

means for processing the received terminal status reports to produce polling response data; and

a memory that stores the polling response data.

260. The processor of claim 259, wherein the processing means comprises:

a demodulator that demodulates the received terminal status reports;

means for reading information fields in the demodulated terminal data, wherein the information fields are appended to terminal identification number fields;

means for sorting the information fields in the received terminal status reports by terminal identification number; and

a temporary memory that accumulates the sorted information fields for the terminals, wherein the accumulated sorted information fields produce the polling response data.

261. The processor of claim 259 or 260, further comprising means for updating the programs watched data in the database with the polling response data.

262. The processor of claim 259, 260 or 261, wherein the counting means comprises:

means for reading the programs watched counts, wherein a separate programs watched count is assigned to different time slots in a day for each program category;

means for sorting the programs watched counts from highest to lowest; and

means for matching the sorted programs watched counts with the programs categories, wherein counts of the programs categories are thereby produced.

263. The processor of claim 262, further comprising:

means for sorting the stored programs watched information by a plurality of program categories;

means for ranking the plurality of program categories by frequency of programs watched in each category, wherein program categories with more programs watched are ranked higher than program categories with less programs watched;

means for selecting a set of highest ranked targeted program categories;

means for choosing individual programs from the set of highest ranked programs categories;

means for packaging the chosen individual programs for transmission to the terminals for display, wherein packages of programs are produced; and

means for transmitting the packages of programs to the terminals over the cable television delivery system.

264. The processor of claim 263, wherein the terminals have the capability of receiving polling request messages and transmitting polling responses upstream over one of a cable television a cable distribution

network in the cable television program delivery system and a telephone line, the polling responses including the programs watched information.

265. A method for accumulating data in a cable television delivery system from a plurality of set top terminals in a cable television delivery system, comprising:

gathering programs watched data for the set top terminals;

storing the gathered programs watched data in a database;

accessing the stored programs watched data;

counting the accessed programs watched data to determine a frequency of programs watched at the set top terminals, wherein programs watched counts are arranged in a programs watched matrix.

266. A method of gathering programs watched information from a set top terminal that acquires programs from a programming source, comprising:

receiving a plurality of programs at the set top terminal;

selecting a program at the set top terminal from the plurality of programs;

monitoring the program selected at the set top terminal;

generating programs watched information relating to the selected program; and

storing the generated programs watched information.

267. The method of claim 266, wherein the steps of monitoring, generating, and storing occur in a card that receives signals from the set top terminal, and wherein the method further comprises the step of transferring signals from the set top terminal to the card.

268. The method of claim 266 or 267, wherein the steps of monitoring, generating and storing occur in a network controller coupled to the set

top terminals, and wherein the method further comprises the step of transferring signals from the set top terminal to the network controller.

269. The method of claim 268, wherein the network controller is located at a cable television headend.

270. The method of claim 268 or 269, wherein the network controller is located at an operations center.

271. The method of claim 268, 269 270, further comprising:

accessing the stored programs watched data; and

arranging the accessed programs watched data in at least one programs watched matrix, wherein one of the at least one programs watched matrices includes time slots and program categories.

272. A television system that receives and stores data related to television viewing habits at a set top terminal, comprising:

a receiver operably connected to the set top terminal, the receiver receiving programs watched information from the set top terminal;

a processor coupled to the receiver, the processor processing the programs watched information to generate a programs watched matrix; and

a database coupled to the processor that stores the programs watched matrix, wherein the processor is capable of intelligent selection of programs to be sent to the set top terminal.

273. The system of claim 272, wherein the received and stored data includes demographic data.

274. The system of claim 272 or 273, wherein the programs watched information is maintained in programs watched matrix by time and

program category such that a most frequently watched program category can be determined for a given time.

275. The system of claim 274, further comprising a database program, wherein programs are stored for intelligent selection based on program category.

276. The system of claim 275, wherein the program database includes an advertisement database, wherein the programs include advertisements, and wherein the advertisements are stored in the advertisement database for intelligent selection based on program category.

277. The system of claim 275, wherein the database includes a television program database, and wherein the programs include the television programs stored in television program database for intelligent selection based on program categories.

278. The system of any one of claims 275 to 277, wherein programs watched information is maintained for groups of set top terminals, the groups including one or more set top terminals so that a most frequently watched program category can be determined for a given time slot for the groups, and wherein the programs watched information is maintained in a program watched matrix by time and program category.

279. The system of claim 273, wherein the processor accesses the demographic data that is provided by the set top terminal, stores the accessed demographic data in a viewer profile database, and weights the demographic data for the set top terminal, wherein the weights may be used in producing the determined program category.

280. The system of claim 273, wherein the processor accesses group demographic data that is maintained for a group of set top terminals, the groups including one or more set top terminals, wherein the group demographic data is maintained in a group profile database, and wherein the controller weights the group demographic data, the weights being used in producing the determined program category.

281. A system for providing programming in a telecommunications network, comprising:

a terminal that gathers program data from a subscriber;

a first processor, operably connected to the terminal, the first processor analyzing the gathered program data to determine the frequency of programs watched by the subscriber and correlating the analyzed program data with categories of programs, wherein each program category includes at least one program;

a second processor that selects a program from the correlated program categories; and

a transmitter coupled to the second processor and the terminal, the transmitter transmitting the selected program to the terminal.

282. The system of claim 281, further comprising a display that displays the transmitted program.

283. The system of claim 282, wherein the display displays a menu of the selected programs from which a subscriber chooses a program to watch.

284. The system of any one of claims 281 to 283, wherein the terminal and the first processor are located in the set top box coupled to a television.

285. The system of any one of claims 281 to 284, wherein the terminal and the first processor are located at a site remote from the set top terminal.

286. The system of claim 285, wherein the site is a cable television headend, and wherein the telecommunications network is a cable television program delivery system.

287. The system of claim 285, wherein the site is an operations center, and wherein the telecommunications network is a cable television delivery system.

288. The system of any one of claims 281 to 287, wherein the program data is programs watched data and wherein the first processor arranges the programs watched data in at least one programs watched matrix and counts the programs watched data to determine the frequency of programs watched.

289. The system of claim 288, wherein the programs watched matrix is arranged by program category and time slot and wherein the first processor reads the programs watched matrix to determine programs watched counts and sorts the programs watched counts within a time slot from highest to lowest counts.

290. The system of any one of claims 281 to 289, wherein the transmitter transmits the selected programs on a television channel, and the terminal is capable of changing television channels.

291. The system of any one of claims 281 to 289, wherein the first processor selects a plurality of programs, and the transmitter transmits the selected programs on a single television channel, and wherein the terminal is capable of masking undesired programs.

292. The system of any one of claims 281 to 291, wherein the program categories include advertisement categories, each advertisement category including at least one advertisement, and wherein the second processor selects a plurality of advertisements from the advertisement categories, and the transmitter transmits the selected advertisements on a single television channel.

293. The system of claim 292, wherein at least one advertisement within an advertisement category is an infomercial.

294. The system of claim 292, wherein at least one advertisement within an advertisement category is a promotion.

295. The system of claim 291, wherein the second processor polls the terminals to monitor the current program being watched for the occurrence of commercial breaks, and wherein the transmitter transmits the selected advertisements in response to the occurrence of the commercial breaks determined by the second processor.

296. The system of any one of claims 281 to 295, wherein the terminal comprises a memory that stores programs watched data, the system further comprising a receiver operably coupled to the first processor that receives the stored programs watched data from the terminal.

297. The system of claim 296, wherein the second processor polls the terminal and wherein the receiver receives status reports sent by the terminal, the receiver further comprising:

a second memory that stores the programs watched data;

a message generator that generates a polling request message, the message directing the terminal to initiate transmission of the terminal status reports, the terminal status reports containing programs watched information, wherein the first processor processes the received terminal status reports to produce polling response data and the second memory stores the polling response data.

298. The system of claim 296, wherein the terminal transmits a terminal status report that contains the programs watched data to the receiver using a random access method.

299. The system of claim 298, wherein the random access method further comprises a CDMA/CD protocol.

300. The system of claim 296, wherein the terminal transmits a terminal status report that contains the programs watched data to the receiver using a standard telephone line.

301. The system of claim 297, the first processor further comprising a demodulator that demodulates the received terminal status reports, wherein the first processor:

reads at least one information field in the demodulated terminal data;

appends each information field to a terminal identification number field; and

sorts each information field in the received terminal status reports by terminal identification number, and wherein the second memory accumulates the sorted information fields for each terminal, the first processor using the accumulated sorted information fields to produce the polling response data.

302. The system of claim 201, wherein the first processor updates the programs watched data in the second memory with the polling response data.

303. The system of any one of claims 281 to 302, wherein the first processor gathers demographic data from subscribers and correlates the demographic data with the categories of programs.

304. The system of claim 303, wherein a group of subscribers forms a statistically significant number of subscribers, and the first processor uses both the gathered demographic data and simulated demographic profiles as the demographic data, and wherein the first processor:

gathers demographic data from the statistically significant group of subscribers; and

generates 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.

305. The system of any one of claims 281 to 304, wherein the terminal is located in the set top box and the first processor is located at a site remote from the set top box.

306. The system of claim 305, wherein the site is a cable television headend.

307. The system of claim 306, wherein the site is an operations center.

308. The system of any one of claims 281 to 307, wherein the terminal gather programs watched data from a plurality of subscribers, wherein the first processor processes the programs watched data gathered from the plurality of subscribers, and wherein the transmitter transmits the selected programs to the plurality of subscribers.

309. A method for determining program ratings at an operations center using set top terminals, comprising:

receiving programs watched information from the set top terminals in an upstream data transmission at a cable television headend;

transmitting the programs watched information from the cable television headend to the operations center;

gathering the transmitted programs watched information from the set top terminals to form a database of program watched information at the operations center; and

calculating program ratings, wherein the program ratings are calculated using the programs watched database.

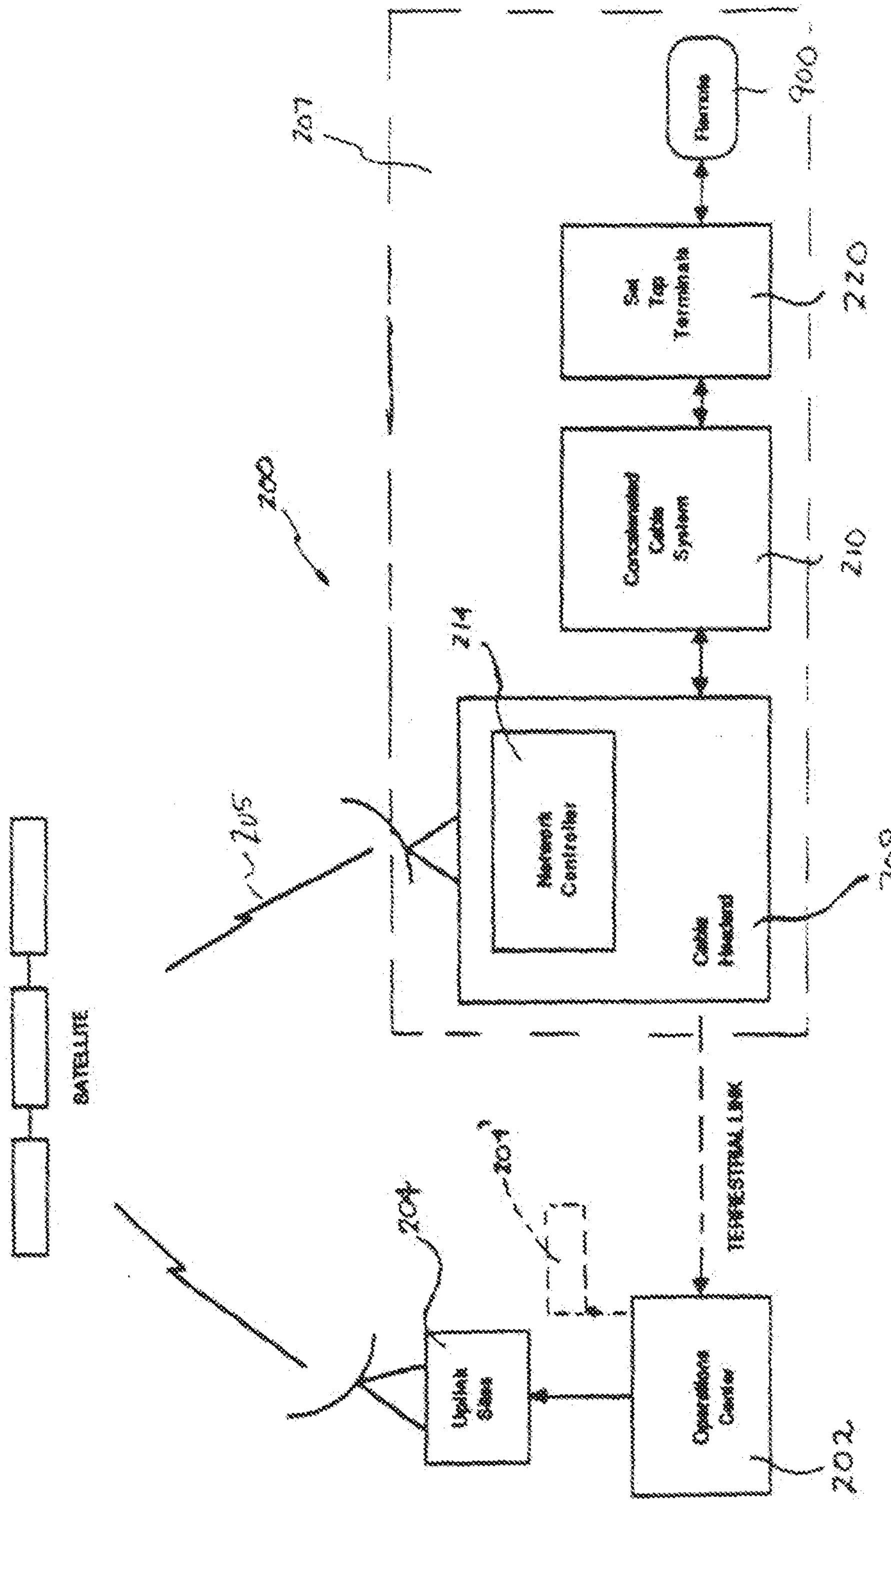
SATELLITE
205
207
200
Network Controller
Cable Headend
Concatenated Cable System
Set Top Terminals
Remote
900
220
210
214
208
TERRESTRIAL LINK
204
204'
Uplink Sites
Operations Center
202

Figure 1

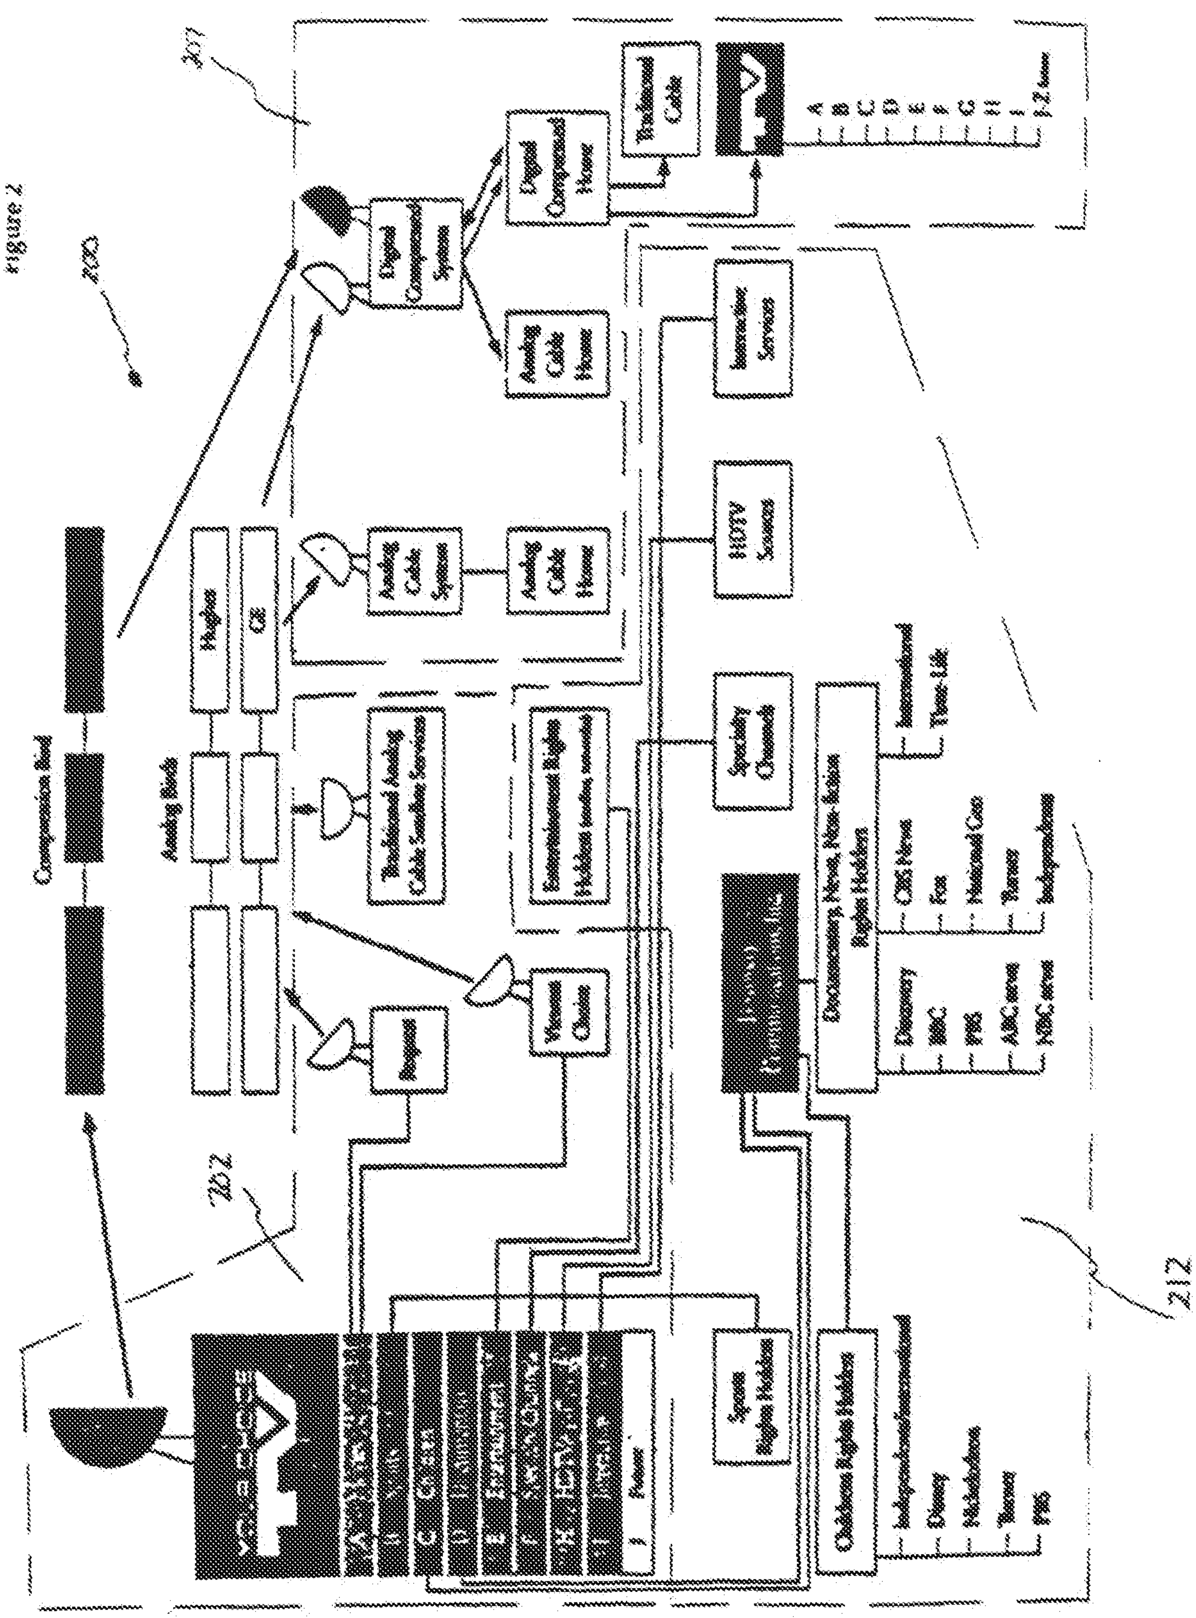
Figure 2
200
202
207
212
Compression Bird
Analog Birds
Hughes
GE
Digital Compression System
Analog Cable System
Analog Cable Homes
Digital Compressed Homes
Traditional Cable
Traditional Analog Cable Satellite Services
Entertainment Rights Holders
Specialty Channels
HDTV Services
Interactive Services
Sports Rights Holders
Childrens Rights Holders
Documentary, News, Non-fiction Rights Holders

Figure 3

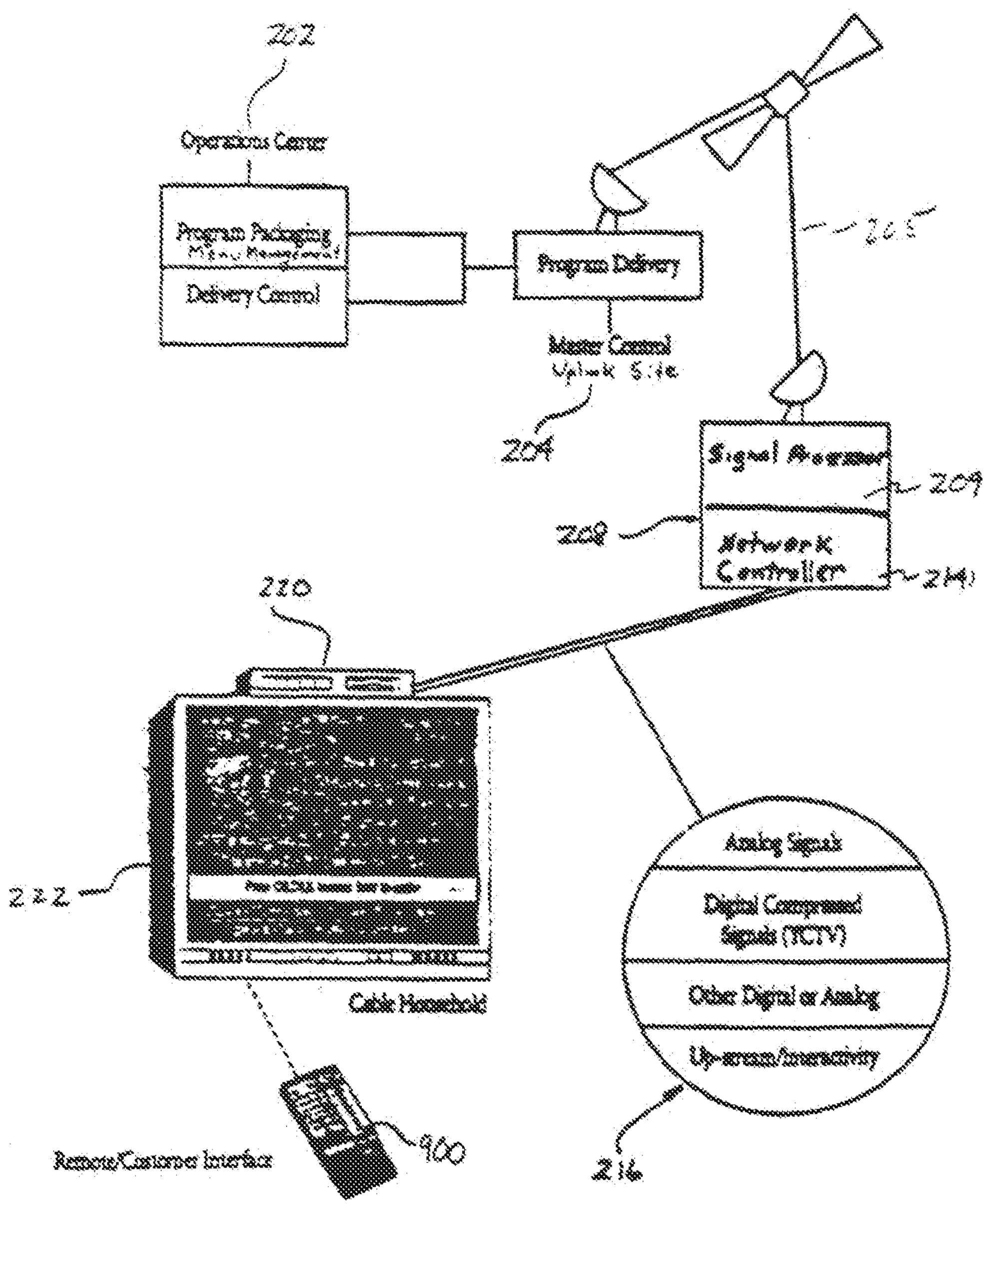
202
Operations Center
Program Packaging
Menu Management
Delivery Control
Program Delivery
Master Control
Uplink Site
204
205
Signal Processor
209
208
Network Controller
214
220
222
Cable Household
Analog Signals
Digital Compressed Signals (TCTV)
Other Digital or Analog
Up-stream/Interactivity
Remote/Customer Interface
900
216

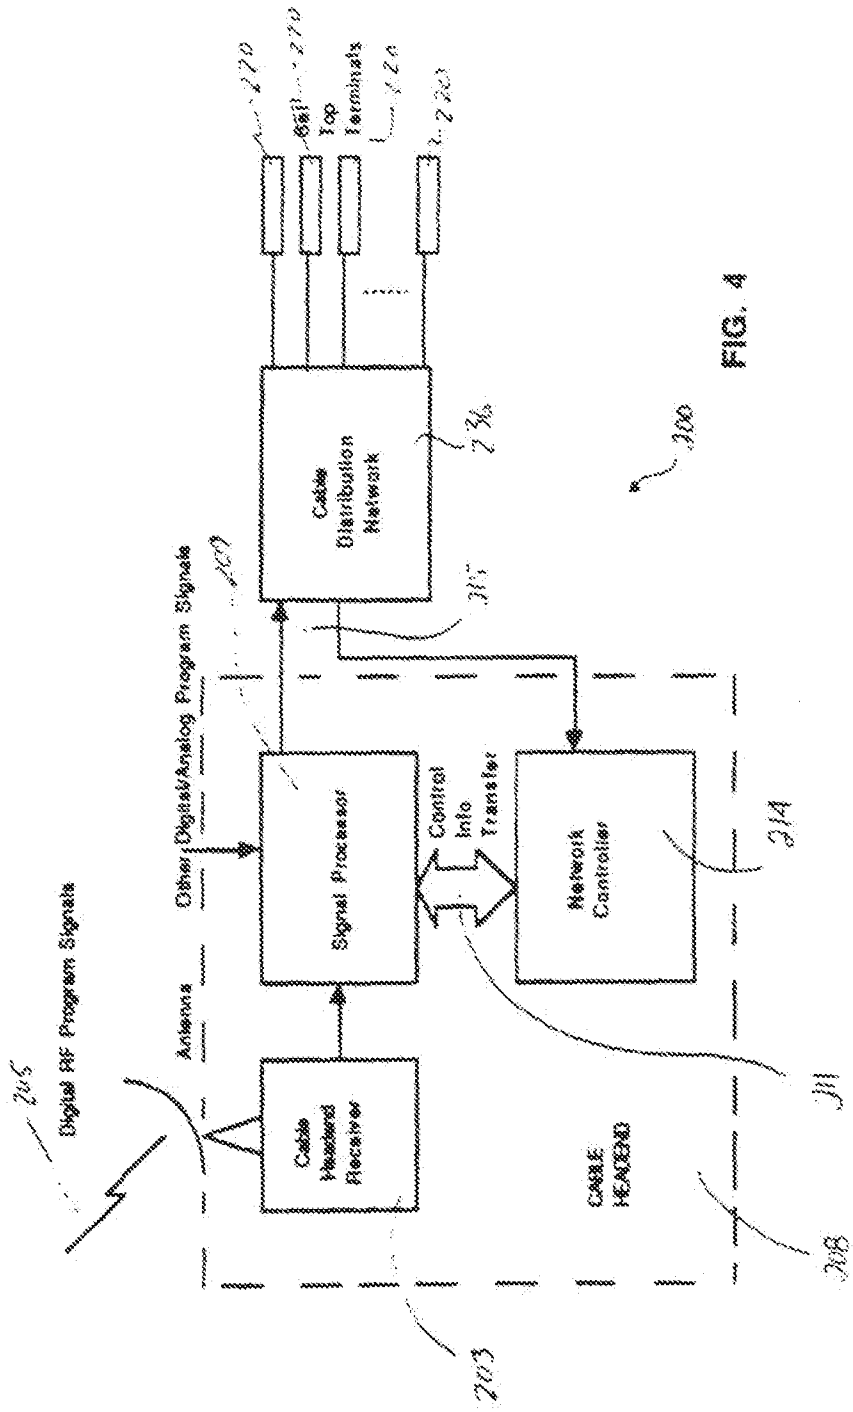
Digital RF Program Signals
205
Antenna
Cable Headend Receiver
203
Signal Processor
209
Other Digital/Analog Program Signals
Control Info Transfer
311
Network Controller
214
CABLE HEADEND
208
Cable Distribution Network
236
315
Set Top Terminals
220
300

FIG. 4

FIG. 5

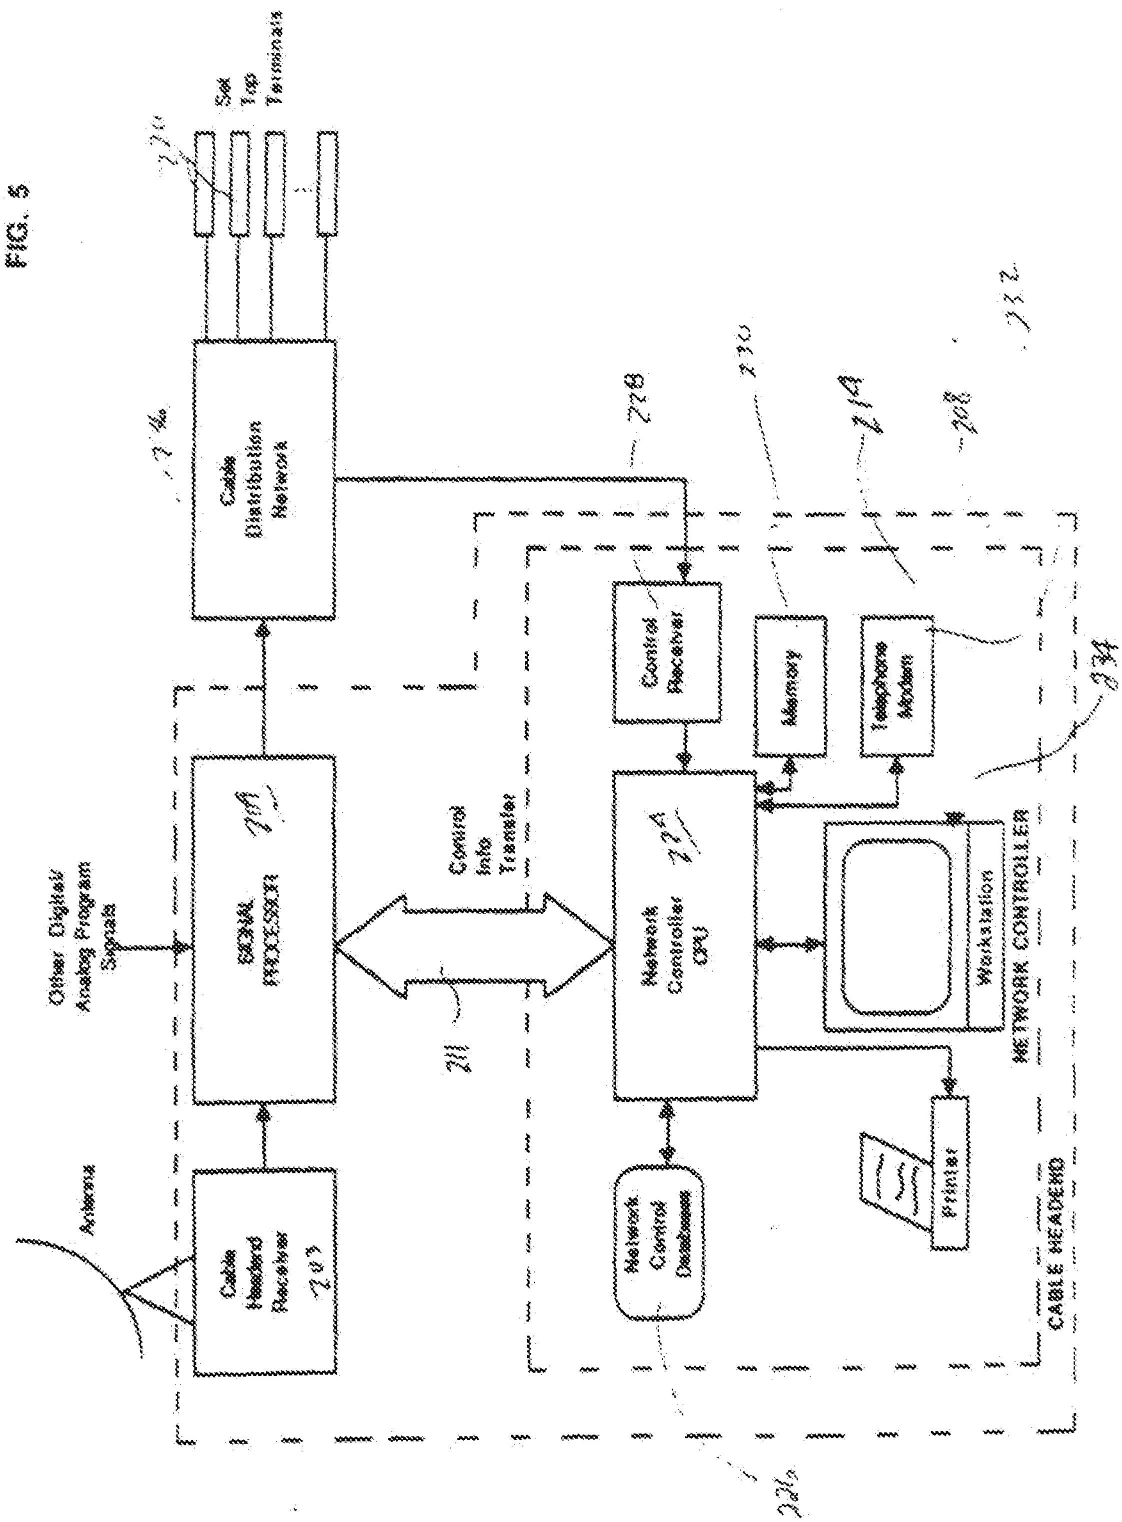
Set Top Terminals
Cable Distribution Network
Other Digital/ Analog Program Signals
SIGNAL PROCESSOR
Antenna
Cable Headend Receiver
Control Info Transfer
Control Receiver
Memory
Telephone Modem
Network Controller CPU
Workstation
NETWORK CONTROLLER
Printer
Network Control Database
CABLE HEADEND

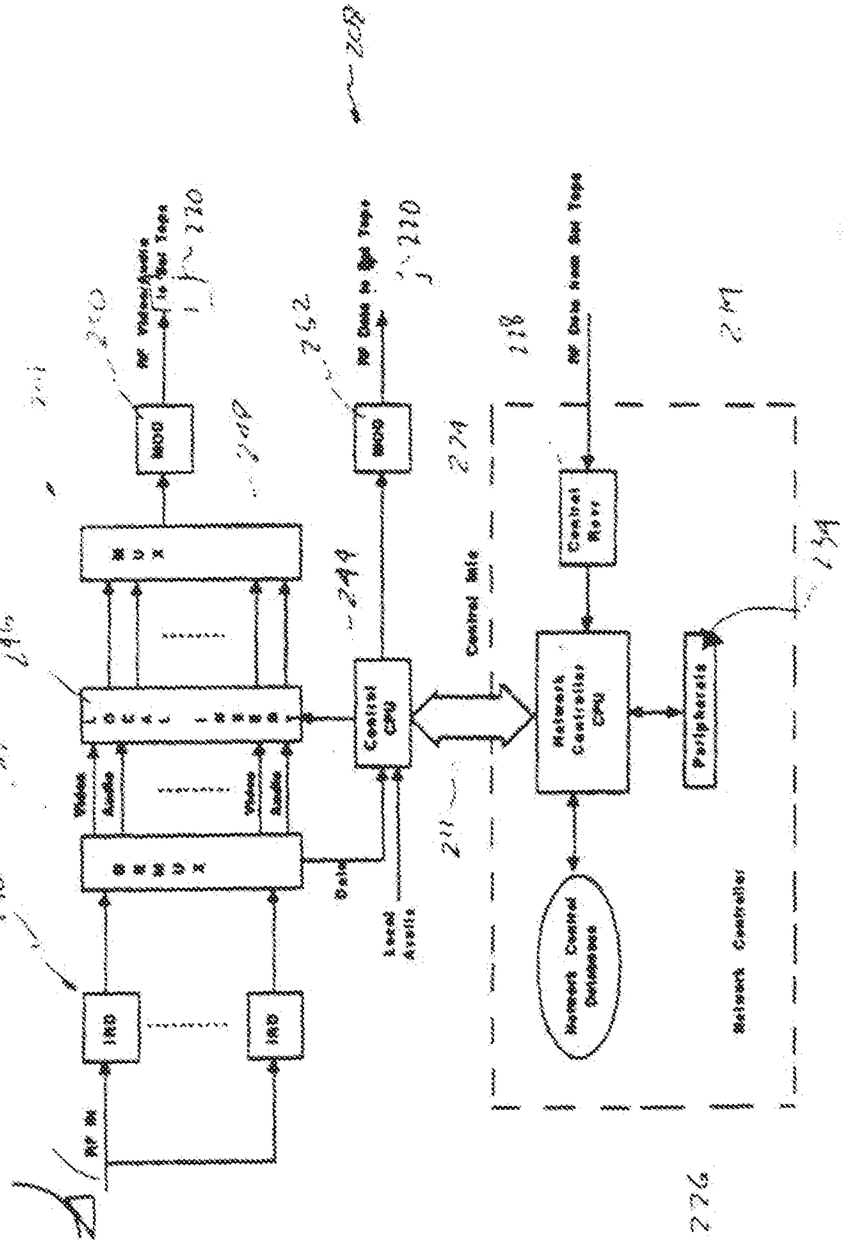

BASIC CABLE HEADEND WITH NETWORK CONTROLLER

Fig. 5a

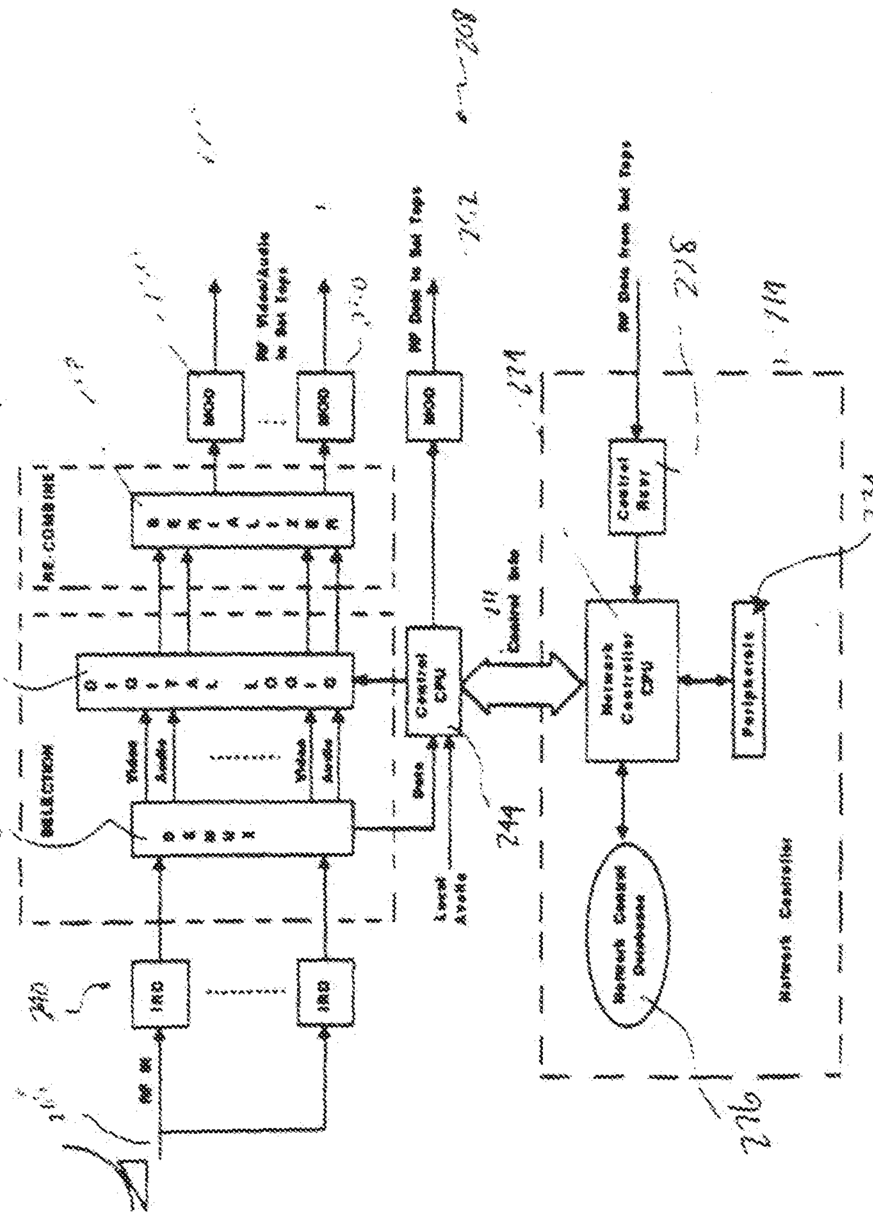

Fig. 6b

CABLE HEADEND WITH COMBINER AND NETWORK CONTROLLER

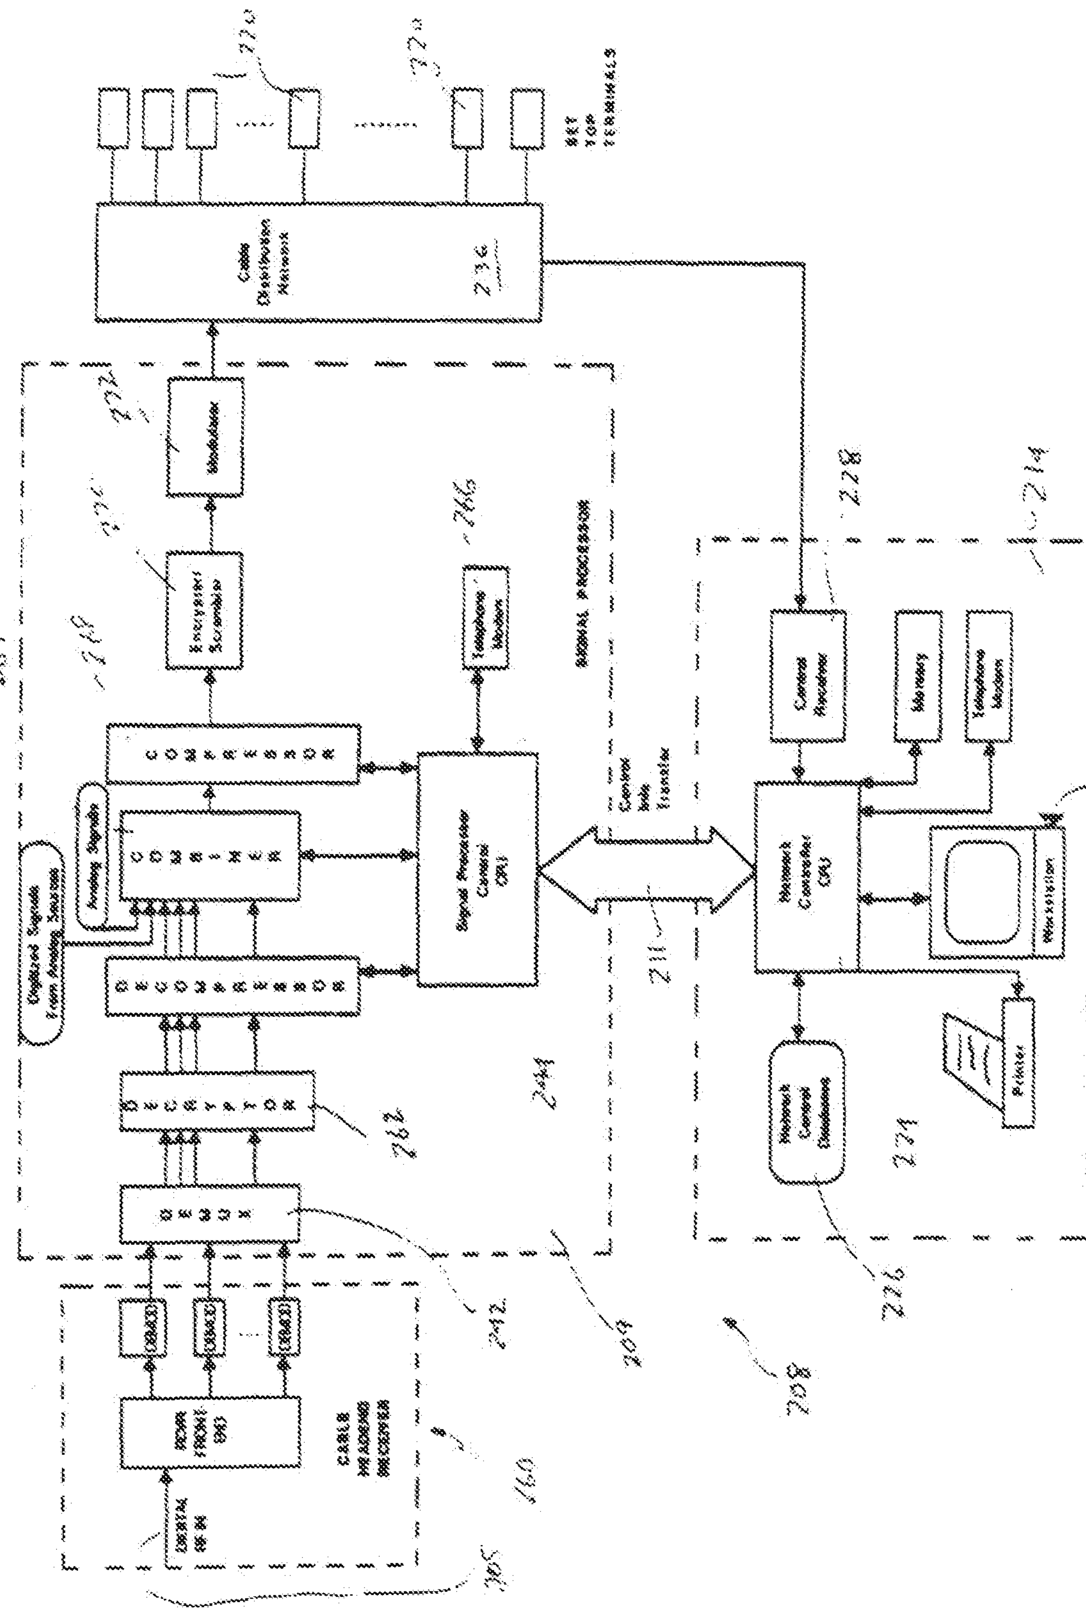
SET TOP TERMINALS
Cable Distribution Network
Modulator
Encryptor Scrambler
Telephone Modem
SIGNAL PROCESSOR
COMPRESSOR
COMBINER
Analog Signals
Digitized Signals From Analog Sources
Signal Processor Control CPU
DECOMPRESSOR
DECRYPTOR
DEMUX
CABLE HEADEND RECEIVER
Control Data Transfer
Control Receiver
Memory
Telephone Modem
Network Controller CPU
Workstation
Network Control Databases
Printer
NETWORK CONTROLLER

FIG. 7

FIG. 8a

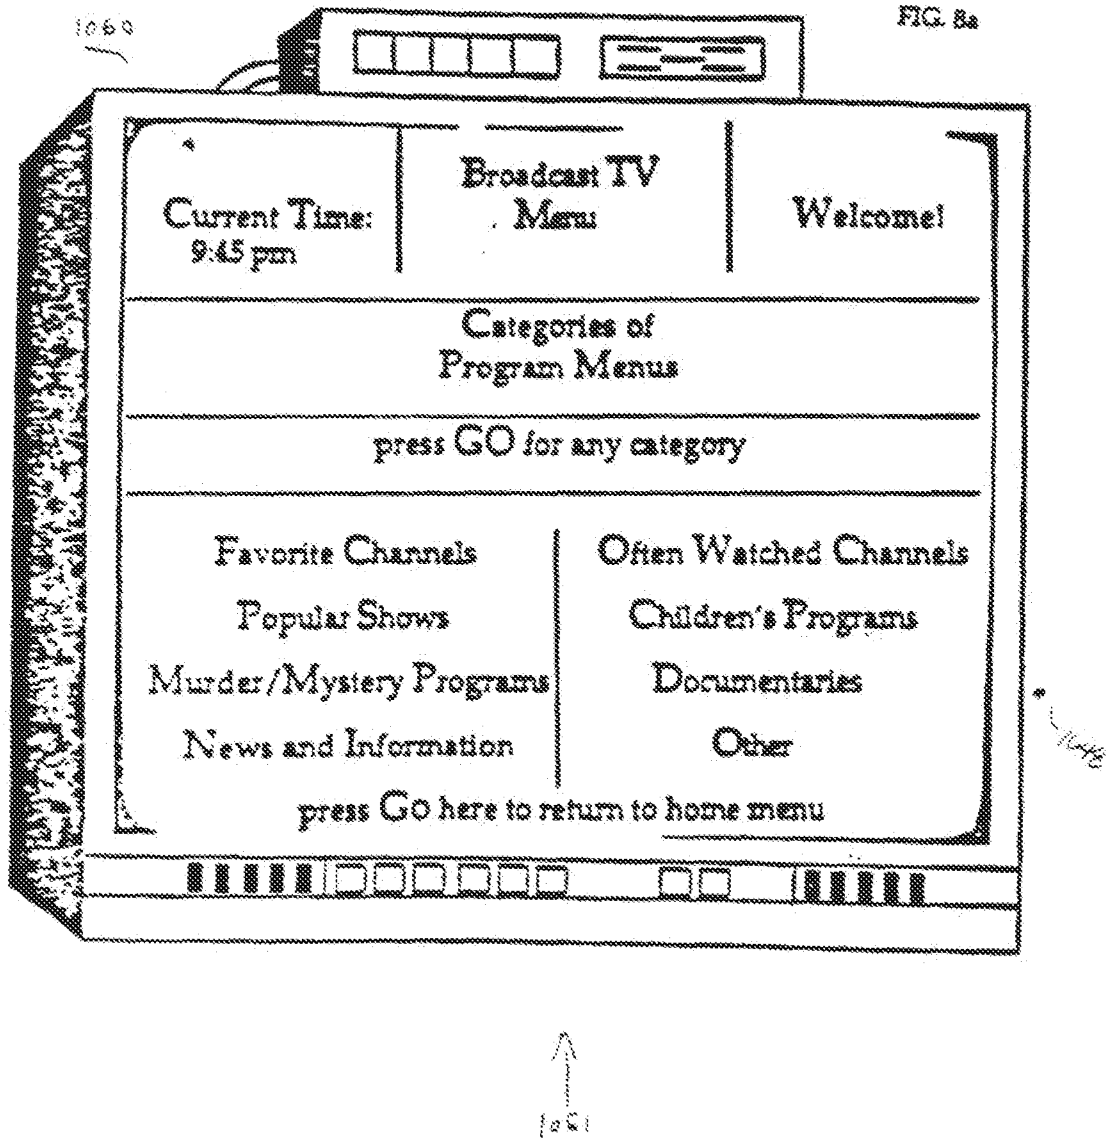
1060
Current Time:
9:45 pm
Broadcast TV
Menu
Welcome!
Categories of
Program Menus
press GO for any category
Favorite Channels
Popular Shows
Murder/Mystery Programs
News and Information
Often Watched Channels
Children's Programs
Documentaries
Other
press Go here to return to home menu
1061

FIG. 8b

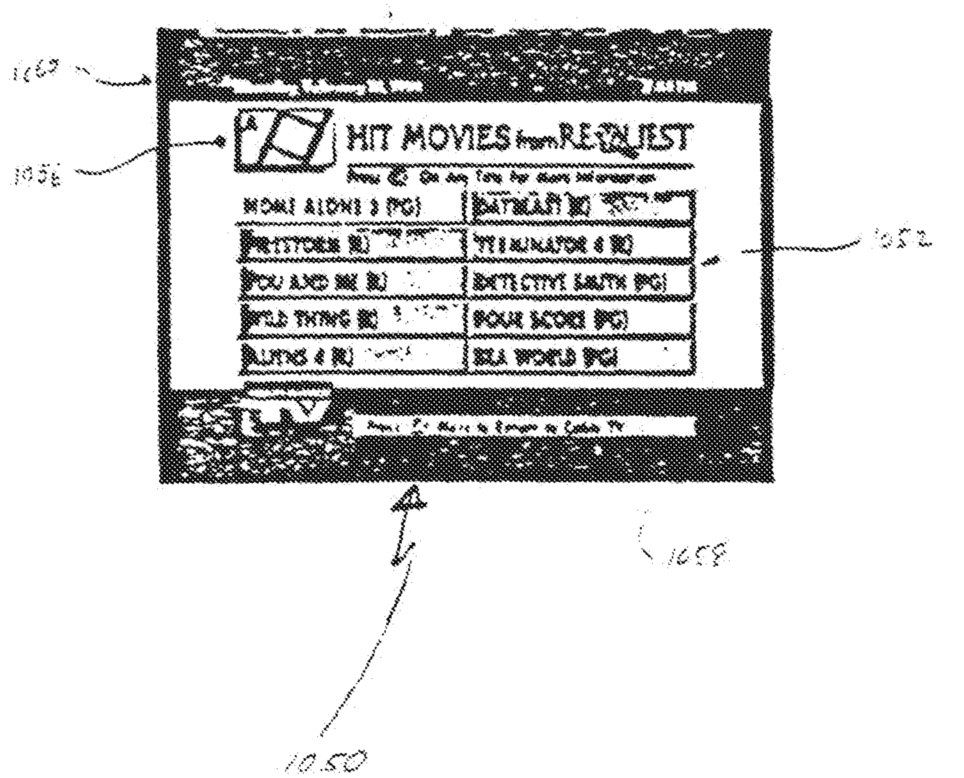

Figure 8c

1060

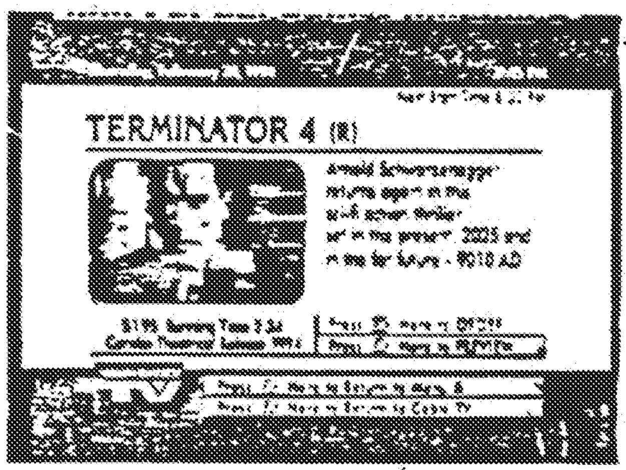

1058

1054

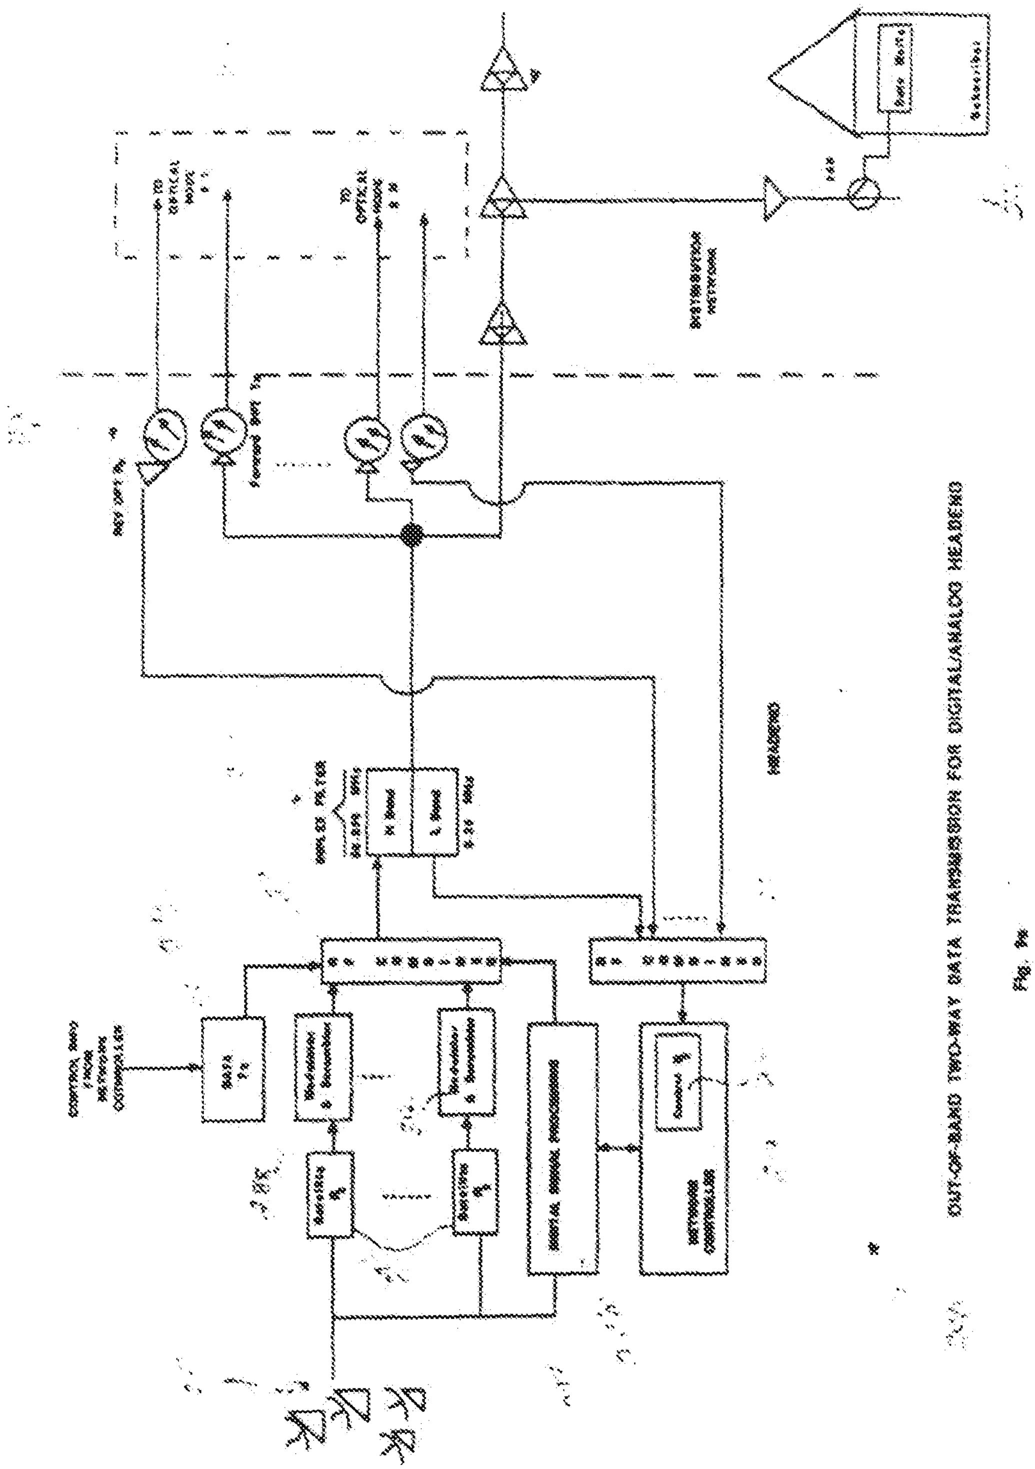

OUT-OF-BAND TWO-WAY DATA TRANSMISSION FOR DIGITAL/ANALOG HEADEND

Fig. 9a

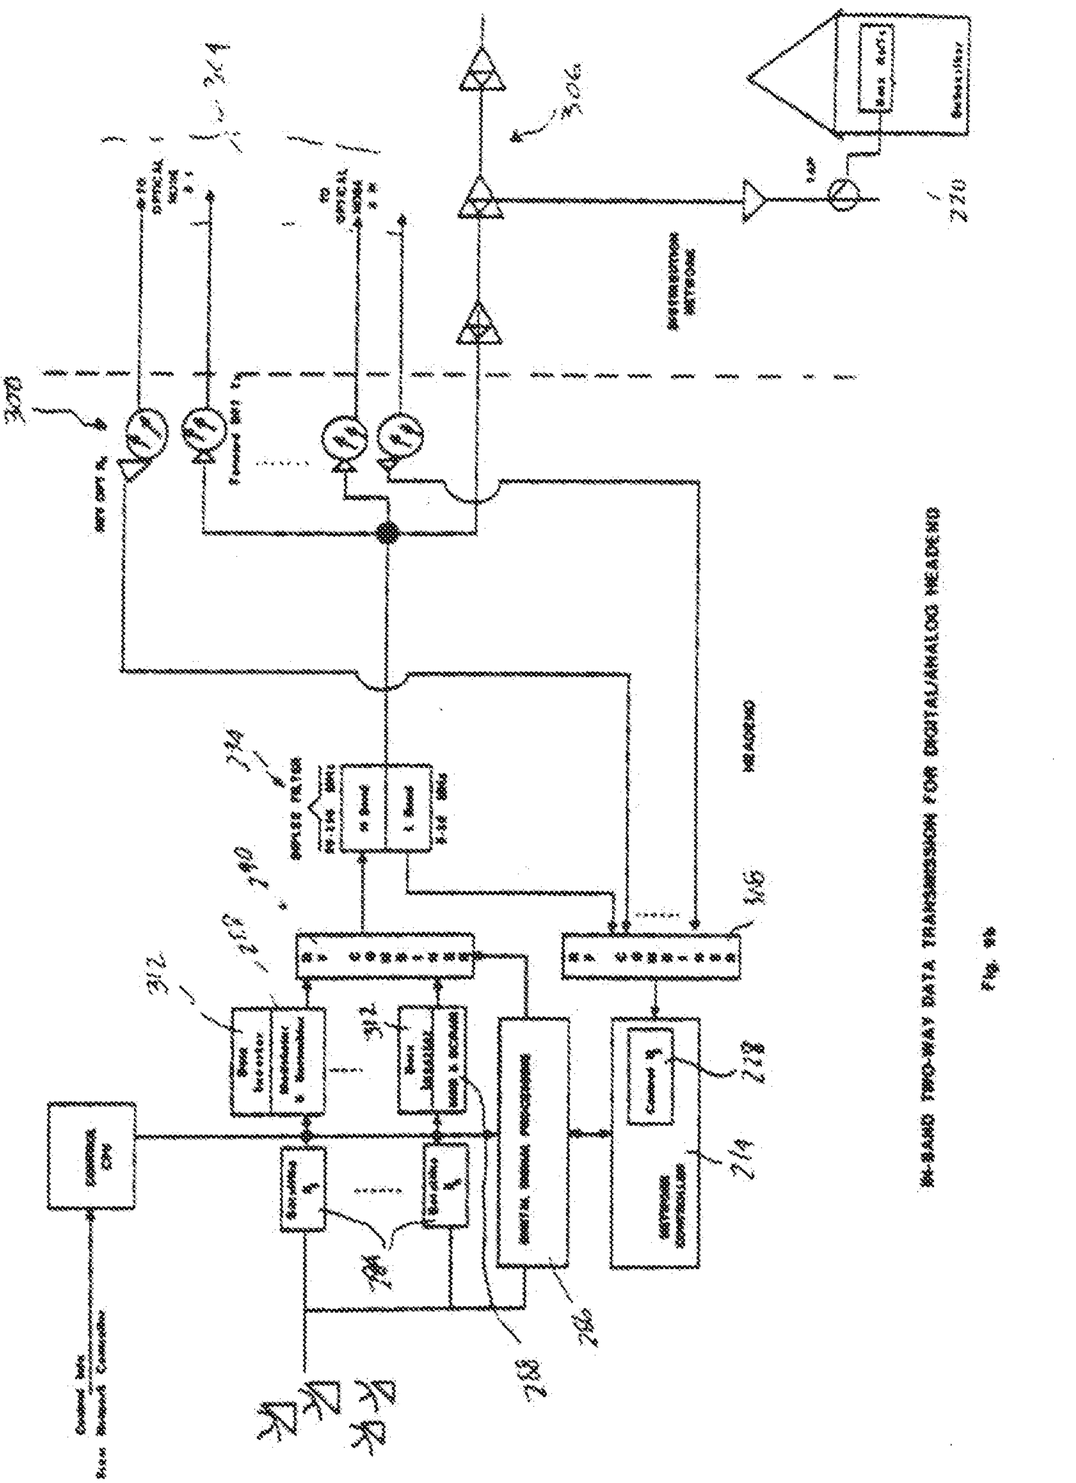

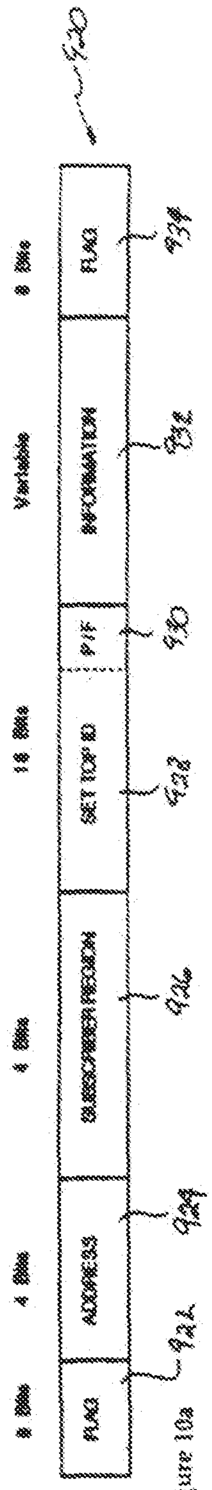

Figure 10a

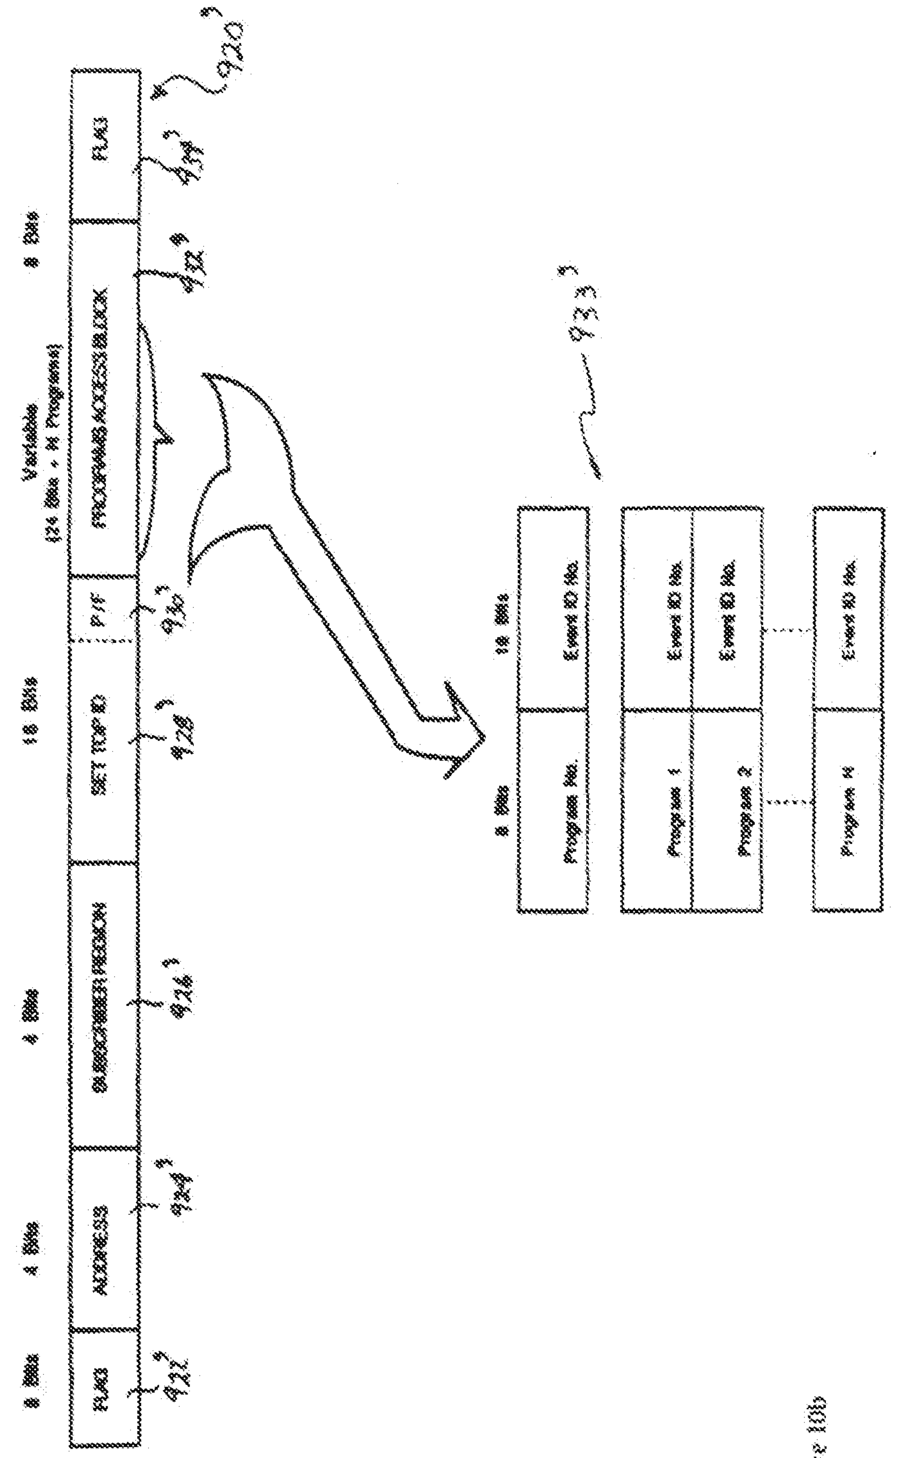

Figure 10b

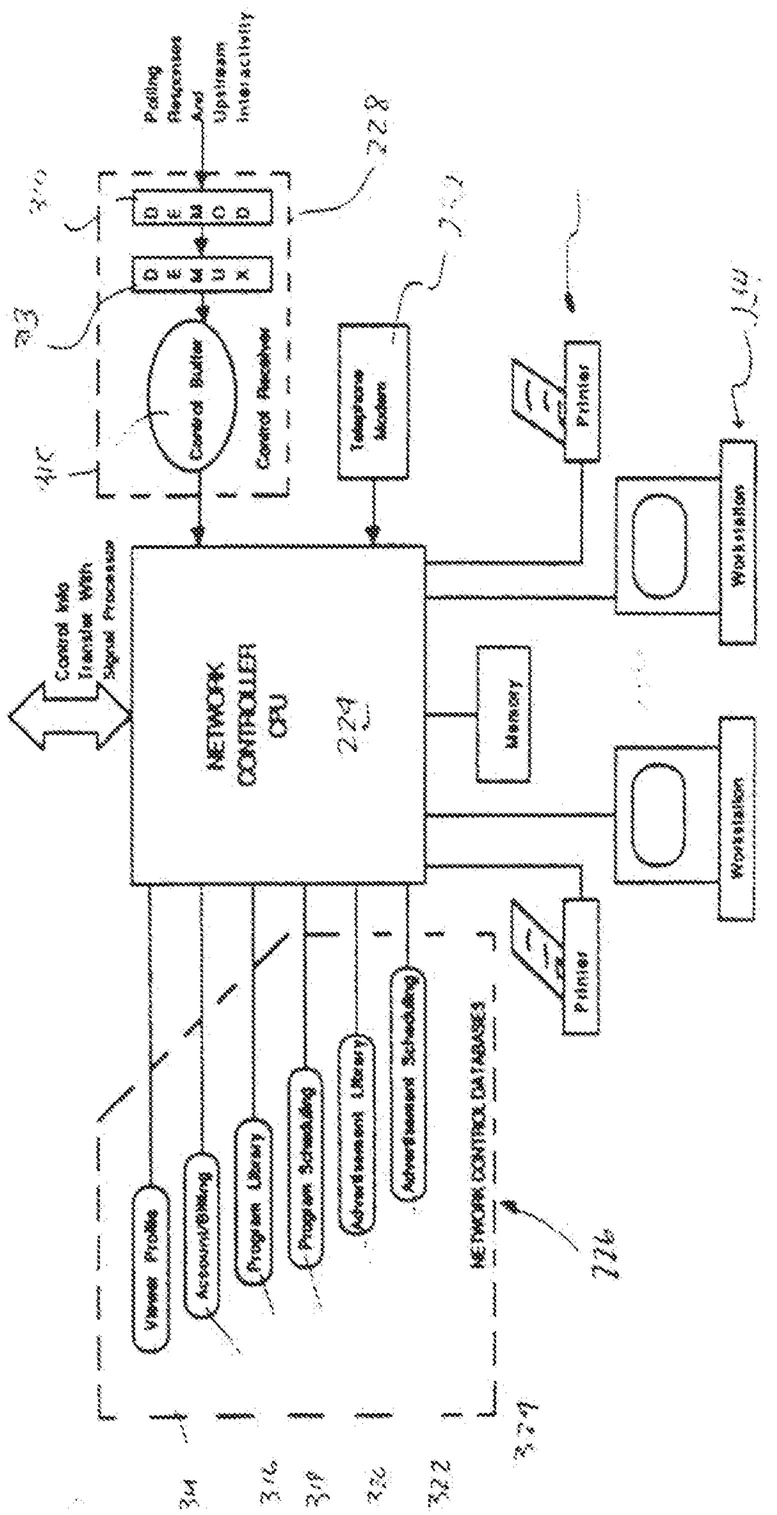
Control Info Transfer With Signal Processor
NETWORK CONTROLLER CPU
224
Polling Responses And Upstream Interactivity
DEMUX
DEMOD
228
Control Buffer
Control Receiver
Telephone Modem
Printer
Workstation
Memory
Workstation
Printer
Viewer Profile
Account/Billing
Program Library
Program Scheduling
Advertisement Library
Advertisement Scheduling
NETWORK CONTROL DATABASES

FIG. 11

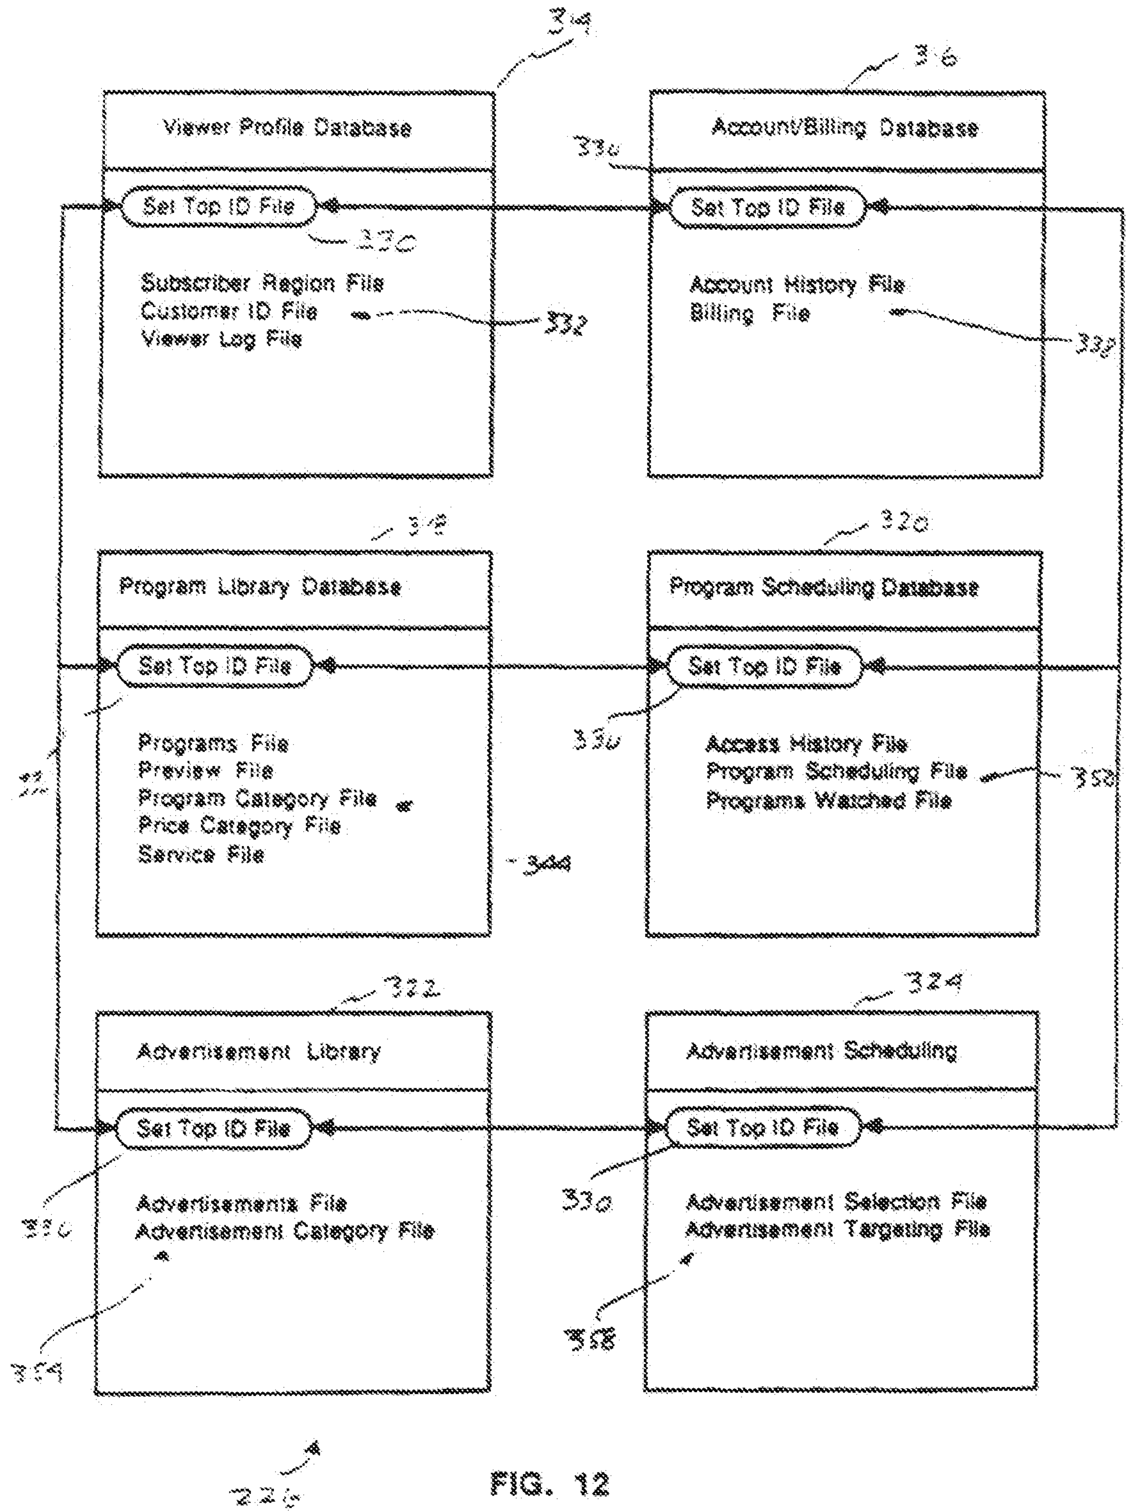
314
316
Viewer Profile Database
Set Top ID File
330
Subscriber Region File
Customer ID File
Viewer Log File
332
Account/Billing Database
330
Set Top ID File
Account History File
Billing File
338
318
Program Library Database
Set Top ID File
Programs File
Preview File
Program Category File
Price Category File
Service File
344
320
Program Scheduling Database
Set Top ID File
330
Access History File
Program Scheduling File
Programs Watched File
350
322
Advertisement Library
Set Top ID File
330
Advertisements File
Advertisement Category File
354
324
Advertisement Scheduling
Set Top ID File
330
Advertisement Selection File
Advertisement Targeting File
358
226

FIG. 12

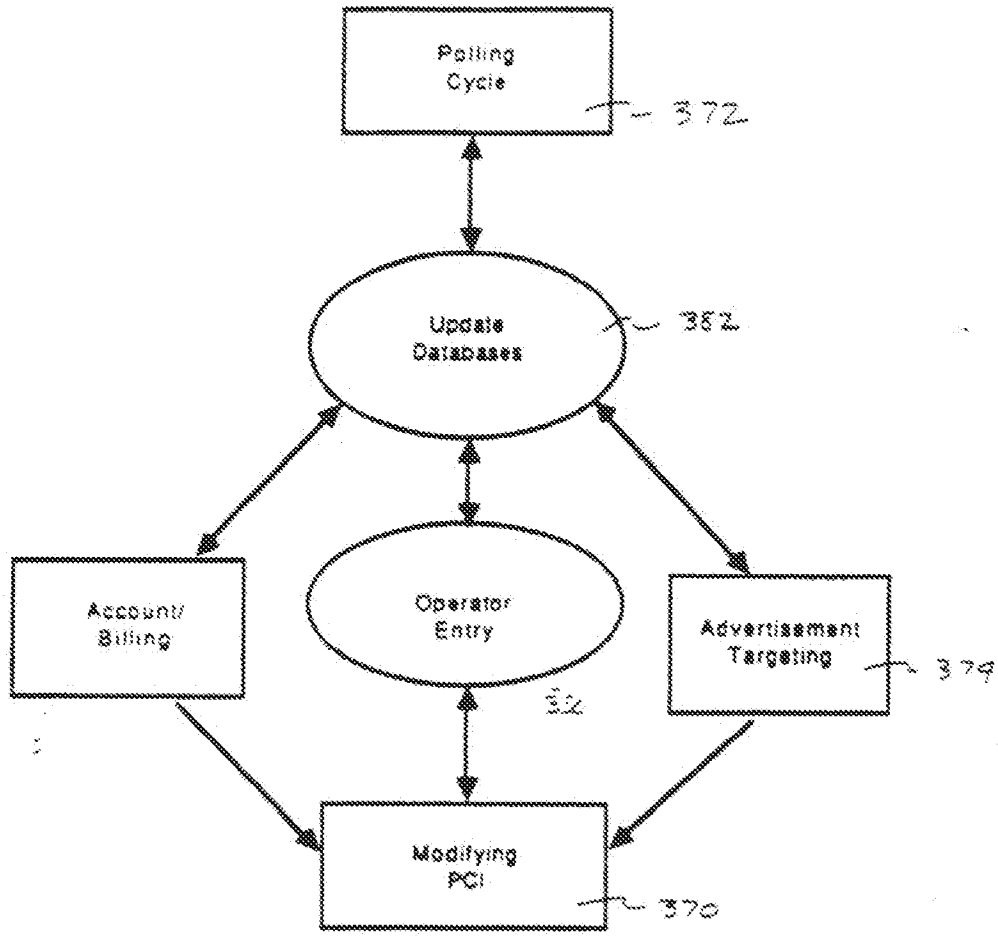
Polling Cycle
372
Update Databases
352
Account/ Billing
Operator Entry
Advertisement Targeting
374
Modifying PCI
370

FIG. 13

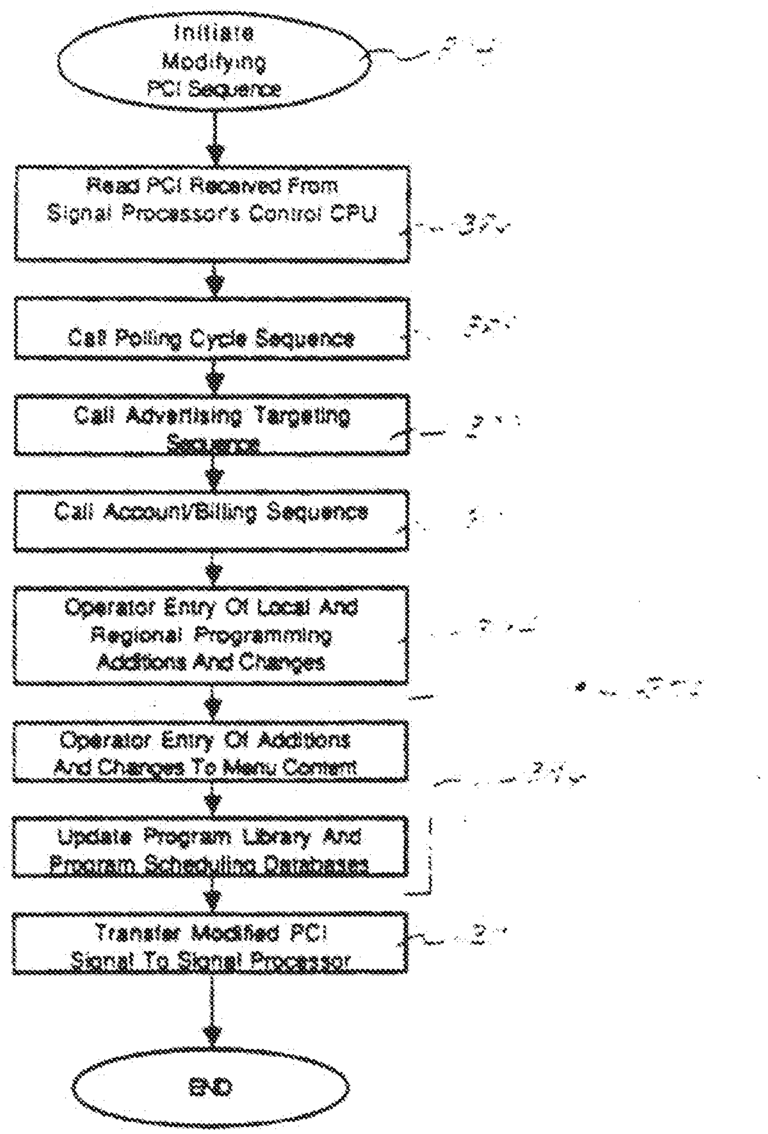
Initiate
Modifying
PCI Sequence
Read PCI Received From
Signal Processor's Control CPU
Call Polling Cycle Sequence
Call Advertising Targeting
Sequence
Call Account/Billing Sequence
Operator Entry Of Local And
Regional Programming
Additions And Changes
Operator Entry Of Additions
And Changes To Menu Content
Update Program Library And
Program Scheduling Databases
Transfer Modified PCI
Signal To Signal Processor
END

FIG. 14

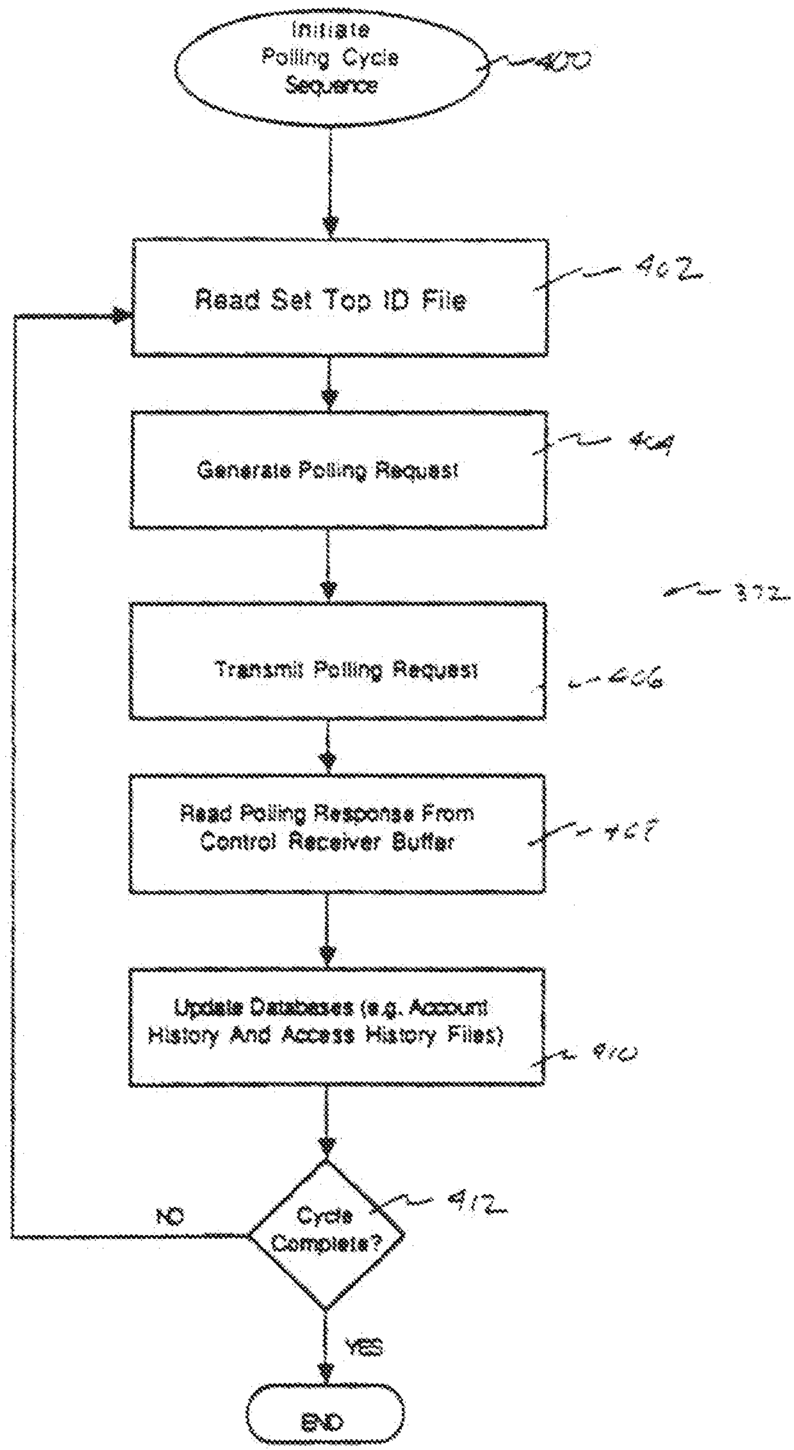
Initiate Polling Cycle Sequence
400
Read Set Top ID File
402
Generate Polling Request
404
372
Transmit Polling Request
406
Read Polling Response From Control Receiver Buffer
408
Update Databases (e.g. Account History And Access History Files)
410
Cycle Complete?
412
NO
YES
END

FIG. 15

MONTHLY PROGRAMS WATCHED MATRIX

351

PROGRAM CATEGORIES

TIME SLOT	SPORTS	NEWS	MOVIES	CHILDREN	ENTERTAINMENT
0000-0400	0	2	2	0	0
0400-0800	0	2	0	3	0
0800-1200	2	0	10	2	6
1200-1600	20	1	3	5	4
1600-2200	8	6	13	0	5
2000-2400	0	10	2	4	2

FIG. 16

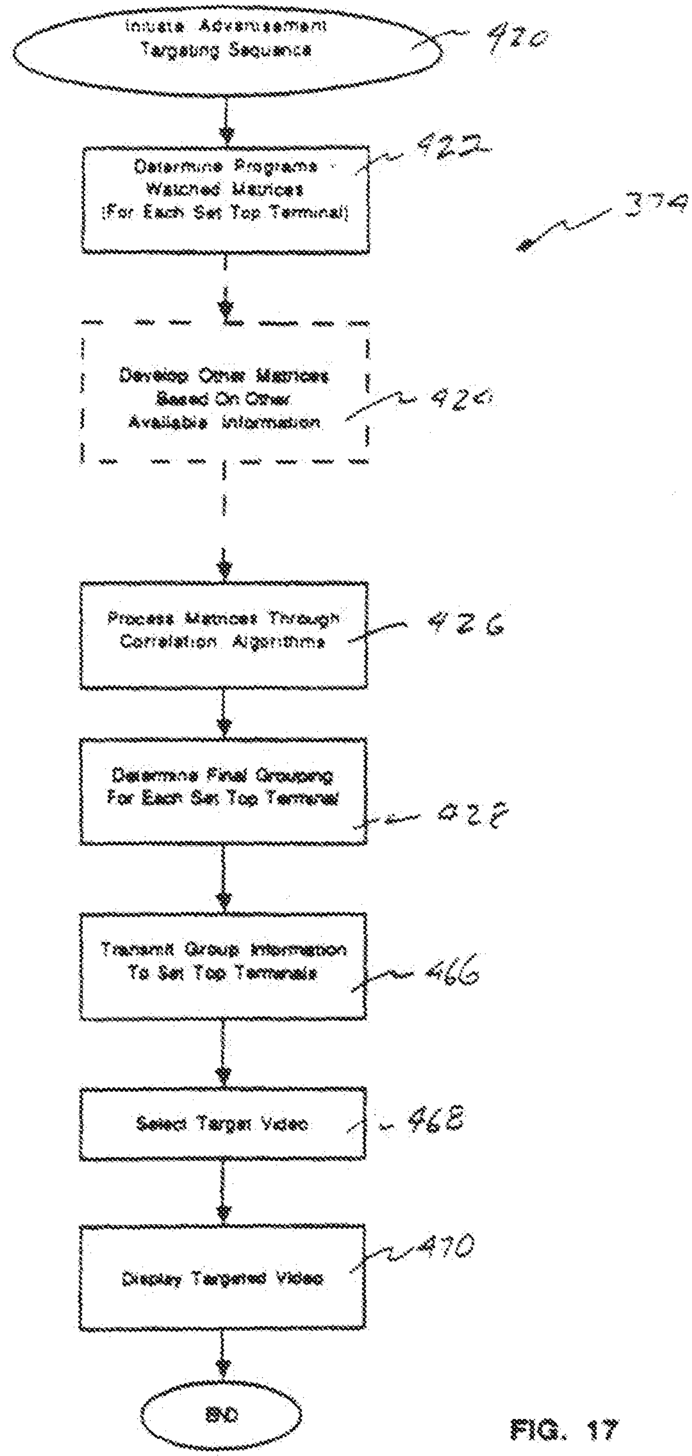
Initiate Advertisement Targeting Sequence
420
Determine Programs Watched Matrices (For Each Set Top Terminal)
422
374
Develop Other Matrices Based On Other Available Information
420
Process Matrices Through Correlation Algorithms
426
Determine Final Grouping For Each Set Top Terminal
428
Transmit Group Information To Set Top Terminals
466
Select Target Video
468
Display Targeted Video
470
END

FIG. 17

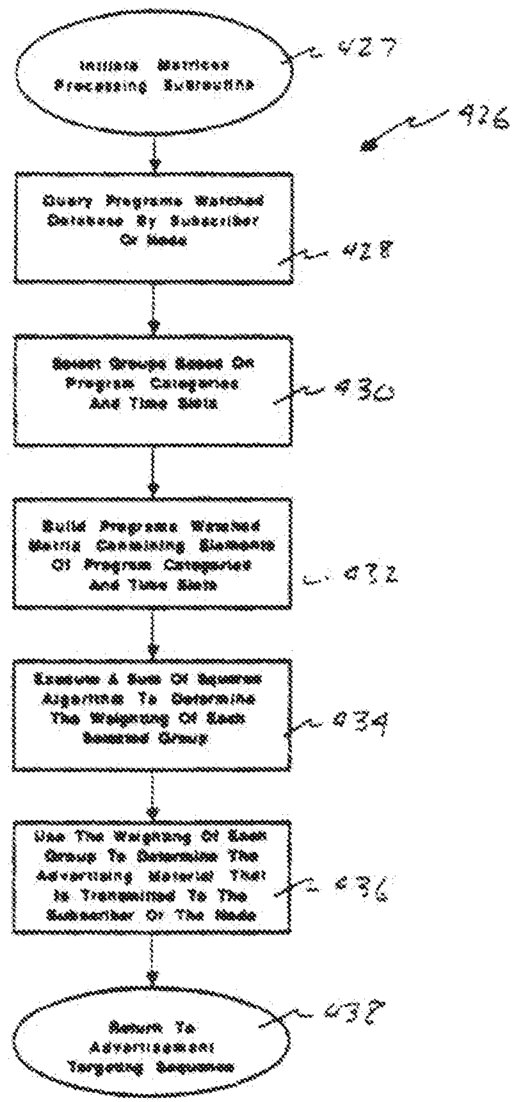
Initiate Matrices Processing Subroutine
427
426
Query Programs Watched Database By Subscriber Or Node
428
Select Groups Based On Program Categories And Time Slots
430
Build Programs Watched Matrix Containing Elements Of Program Categories And Time Slots
432
Execute A Sum Of Squares Algorithm To Determine The Weighting Of Each Selected Group
434
Use The Weighting Of Each Group To Determine The Advertising Material That Is Transmitted To The Subscriber Or The Node
436
Return To Advertisement Targeting Sequence
438

FIG. 18

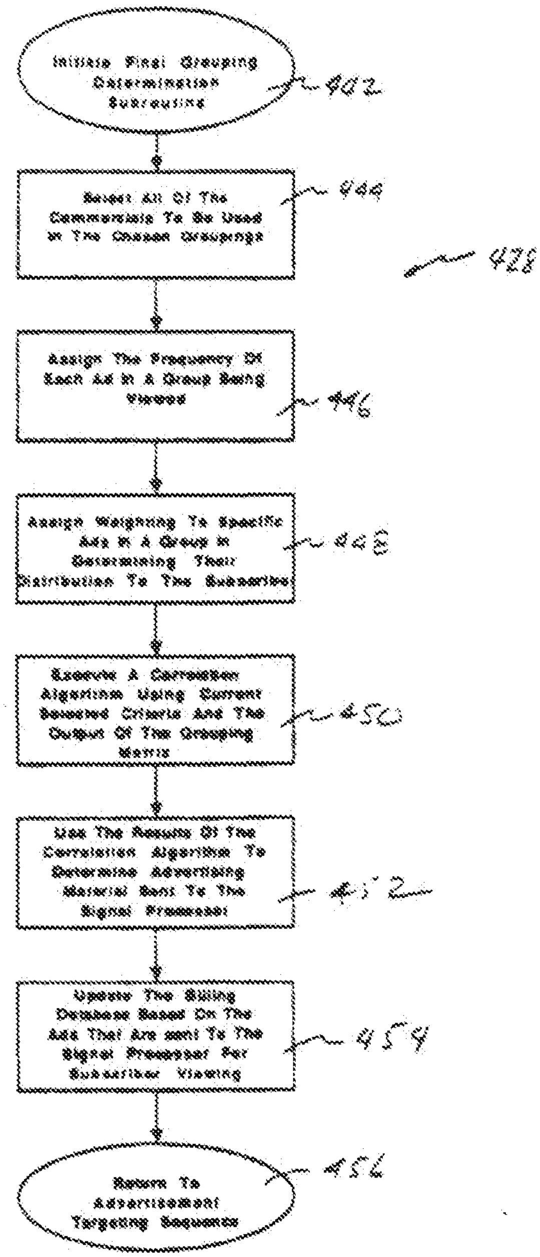
Initiate Final Grouping Determination Subroutine
402
Select All Of The Commercials To Be Used In The Chosen Groupings
444
428
Assign The Frequency Of Each Ad In A Group Being Viewed
446
Assign Weighting To Specific Ads In A Group In Determining Their Distribution To The Subscriber
448
Execute A Correlation Algorithm Using Current Selected Criteria And The Output Of The Grouping Matrix
450
Use The Results Of The Correlation Algorithm To Determine Advertising Material Sent To The Signal Processor
452
Update The Billing Database Based On The Ads That Are sent To The Signal Processor For Subscriber Viewing
454
Return To Advertisement Targeting Sequence
456

FIG. 19

PROGRAM CATEGORIES

160 — Set Top Terminal Groupings; 170; 178

	SPORTS	NEWS	MOVIES	CHILDREN	ENTERTAIN-MENT
A	I	X	III	IV	VI
B	X	II	X	VII	VI
C	I	X	III	D	D
D	IX	VIII	X	IV	V
E	IX	II	X	VII	V

FIG. 20a

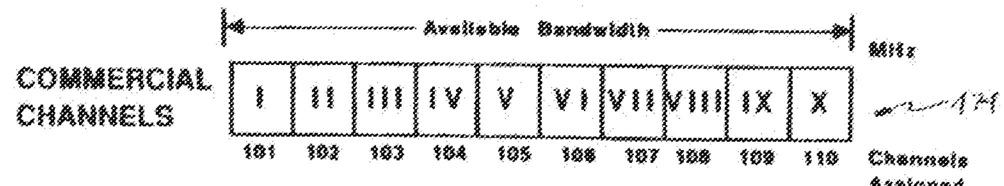

FIG. 20b

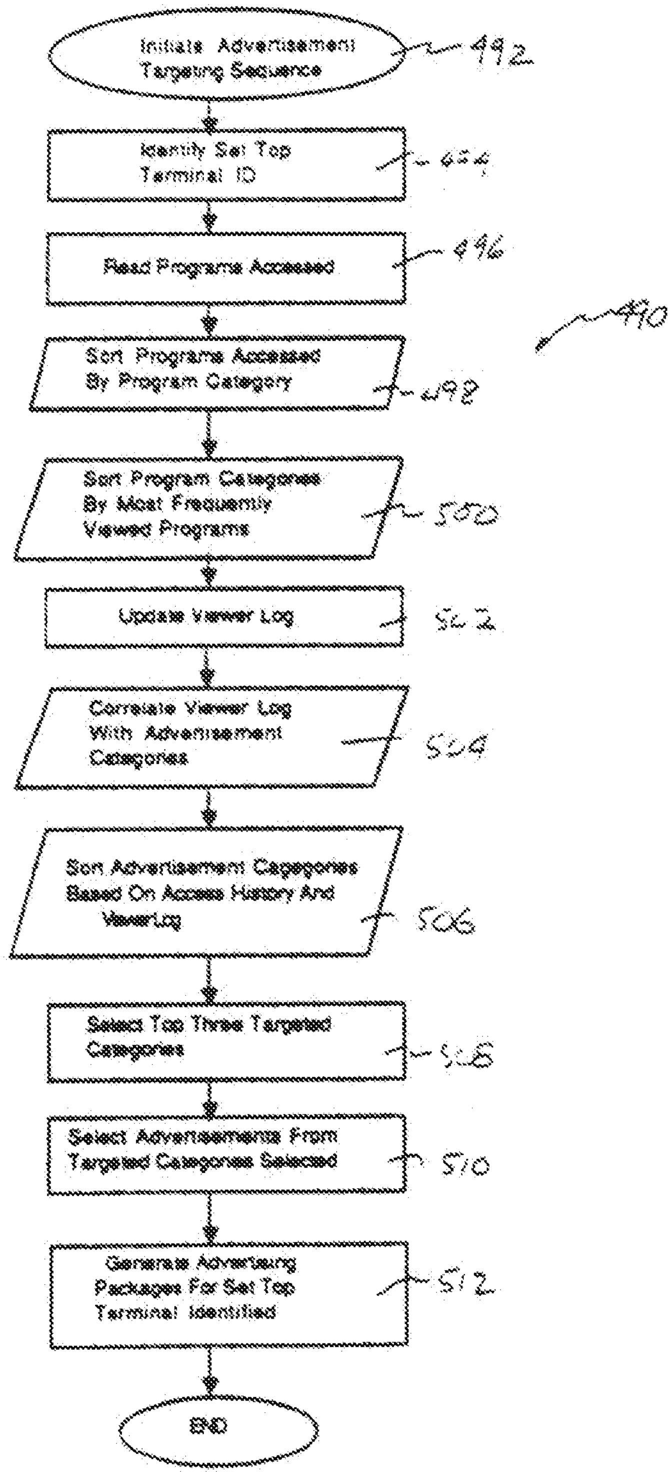
Initiate Advertisement Targeting Sequence
492
Identify Set Top Terminal ID
494
Read Programs Accessed
496
490
Sort Programs Accessed By Program Category
498
Sort Program Categories By Most Frequently Viewed Programs
500
Update Viewer Log
502
Correlate Viewer Log With Advertisement Categories
504
Sort Advertisement Categories Based On Access History And Viewer Log
506
Select Top Three Targeted Categories
508
Select Advertisements From Targeted Categories Selected
510
Generate Advertising Packages For Set Top Terminal Identified
512
END

FIG. 21

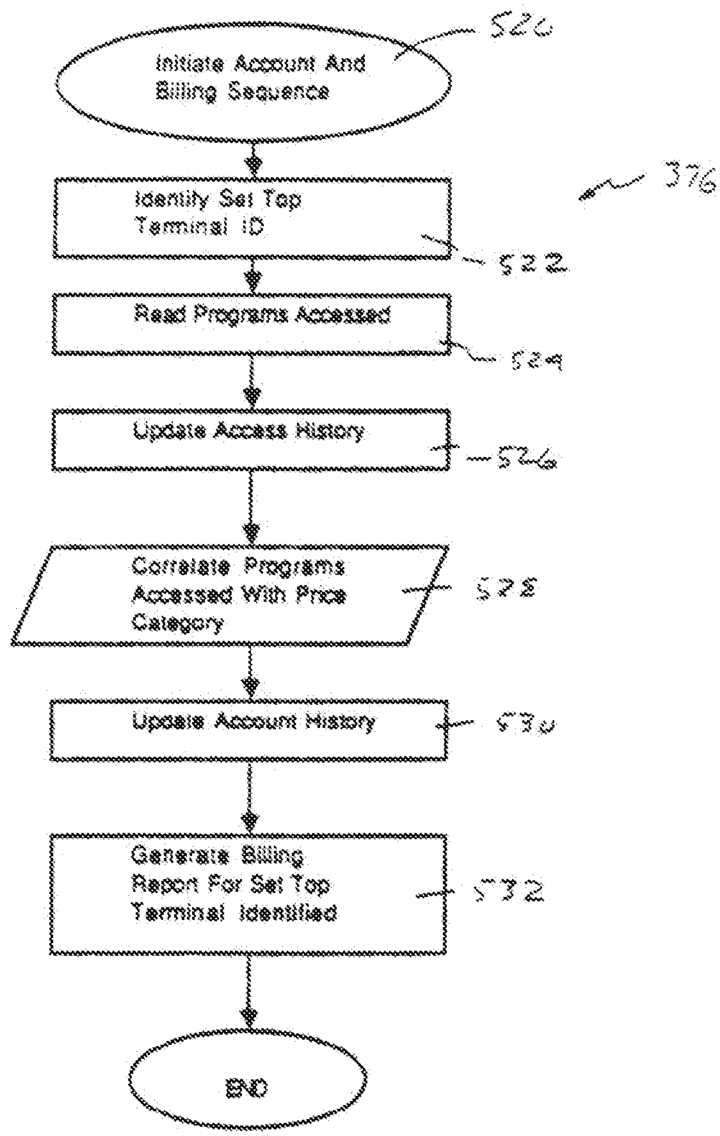
520
Initiate Account And Billing Sequence
376
Identify Set Top Terminal ID
522
Read Programs Accessed
524
Update Access History
526
Correlate Programs Accessed With Price Category
528
Update Account History
530
Generate Billing Report For Set Top Terminal Identified
532
END

FIG. 22

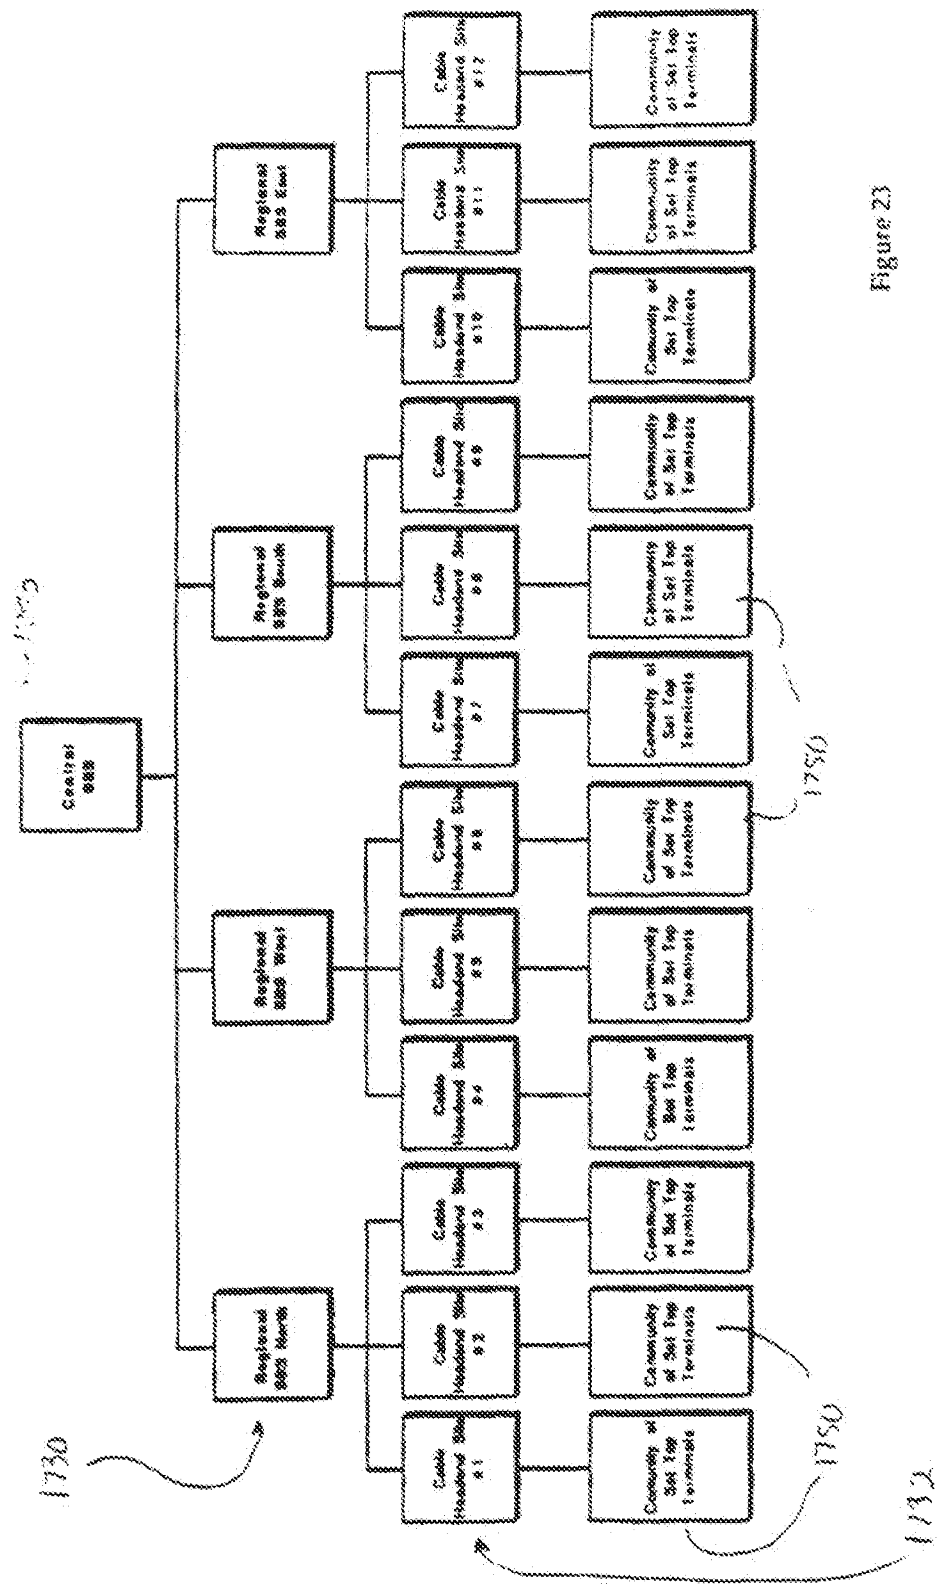

Figure 23

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) Publication number: 0 364 638
A1

(12) EUROPEAN PATENT APPLICATION

(21) Application number: 88309854.3

(22) Date of filing: 20.10.88

(43) Date of publication of application:
25.04.90 Bulletin 90/17

(84) Designated Contracting States:
BE CH DE ES FR GB IT LI NL SE

(51) Int. Cl.5 H04L 12/28

(71) Applicant: International Business Machines Corporation
Old Orchard Road
Armonk, N.Y. 10504(US)

(72) Inventor: Ainscow, Frank
13 Wharf Hill
Winchester Hampshire(GB)

(74) Representative: Bailey, Geoffrey Alan
IBM United Kingdom Limited Intellectual Property Department Hursley Park
Winchester Hampshire SO21 2JN(GB)

(54) Communication network.

(57) In a communication network for communication between a plurality of nodes over a multi-channel communication medium 6, the network comprising a plurality of allocatable channels 30, 30′, 30″, ... for data communication between nodes and a control channel 28 for allocating the allocatable channels to individual nodes, for example node 20, at least one of the channels operates under a token-passing protocol with a limited message packet length whereby selected messages may be sent between nodes with a predetermined maximum delay.

The token-passing channel may be the control channel 28, and the message packets transmitted by the nodes over the control channel may comprise messages for negotiating the allocation of the allocatable channels to the individual nodes. Alternatively, or additionally, message packets comprising data messages unconnected with negotiating the allocation of the allocatable channels may be sent over such a token-passing channel.

More than two nodes may be permitted to communicate on one of the allocatable channels, thereby establishing a logical sub-network which can operate using any desired protocol over the allocated channel.

EP 0 364 638 A1

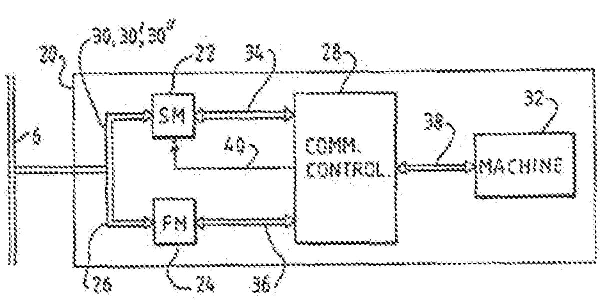

FIG. 3

COMMUNICATION NETWORK

The present invention relates to a communication network for communication between a plurality of nodes over a communication medium.

In order to control the interchange of information over a network in which a plurality of nodes are competing for the use of the network, it is conventional to use some form of network protocol.

An example of a network based on one class of protocol, usually called a token-passing protocol, is a token-bus Local Area Network (token-bus LAN). In a token-passing network each node in the network may only transmit information to other nodes when it is in possession of a "token", sent to it by a predecessor node (logically, not necessarily physically) in the network. A node in possession of the token may pass as many messages as it requires, subject to restraints imposed on it by the protocol under which the network is operating, and then passes the token on to a successor (again logically, not necessarily physically) in the network by sending a message including the token to that node. This node then, similarly, transmits as required and then passes the token on to its successor node. The predecessor and successor relationships are arranged in the network to define a logical ring such that the token is passed repeatedly around this logical ring.

By means of this mechanism, only one node is authorised to transmit at any one time, with the result that collisions, which might otherwise occur if two or more nodes tried to transmit at once, are avoided (barring error conditions). Also, if each node is limited to a maximum token retention time on receipt of the token, then the total token rotation time around the logical ring is guaranteed to be below some maximum figure, this figure being dependent on the maximum token retention time at each node, transmission delays and the number of nodes in the network. Hence, the delay between a node requiring to transmit some information and successfully transmitting that information is a known quantity. This is a significant advantage in, for example, a LAN in a manufacturing environment where one node may be the controller of a machine tool and another a control computer giving the machine instructions.

In order for a communication network to be successful, there are many conflicting requirements, such as compatibility amongst the attached devices, performance, flexibility and reliability. By way of example, consider the situation with regard to industrial LANs. The desire to ensure compatibility, was recognised in 1978 by the International Standards Organisation which proposed and defined the Open Systems Interconnection (OSI) model as a basis for the creation of standards for communication between computing devices. Building on this model, a standard has been proposed for an industrial LAN called the Manufacturing Automation Protocol (MAP). The MAP protocol is described in a book by J. Dwyer and A. Ioannou entitled "MAP and TOP" (ISBN 1 -85091 356 2). The MAP protocol achieves the requirements for compatibility very well, but at some cost to the other requirements as discussed in an article by P Ainscow entitled "Milestones in MAP", published in "Computer Bulletin", Vol 4, Part 2, June 1988, pages 12 and 13, by the British Computer Society.

In pure information handling applications, and in interactions with people, a few seconds delay in the transmission of a message over a communication network is tolerable in most cases. By contrast, where computers are being used to control the operation of mechanical devices, there may be a need for "real time" operation - that is a guaranteed maximum delay time in the order of milliseconds as opposed to hundreds of milliseconds or longer. As indicated above, a token-passing technique of giving access to the communication channel gives the possibility of guaranteeing a maximum response time. The MAP is based on a token-passing protocol. However, the need to handle efficiently long files under the protocol means that the token retention time is relatively high with the result that the time guaranteed is in the order of one second, even for moderately sized networks.

One technique for reducing the guaranteed time is to subdivide a network into a plurality of smaller networks. This approach is not without its difficulties, however, as it introduces the problem of inter-network communication at gateways, increases the transmission delay to nodes on the other networks and means compromising compatibility and flexibility.

For communication networks where a guaranteed maximum response time is not a requirement there exists a class of protocols termed Carrier Sense Multiple Access/Collision Detection (CSMA/CD) protocols. An example of such a protocol is the "Ethernet" protocol. In a CSMA/CD network, any node requiring access may attempt transmission at any time when the network is not active. Collisions between two or more nodes which attempt transmission at the same time are handled by all nodes ceasing transmission then attempting transmission some time later. The delay before attempting retransmission is variable, to avoid two or more nodes repeatedly issuing colliding transmissions, leading to inability to transmit. This type of network has the advantage that at low traffic levels, the delay to transmit a message is generally low. However, it has the disadvantage that at high traffic levels, the delay to transmit a message can get very long indeed, due to frequent collisions occurring, and consequently an unacceptable delay may be experienced in transmission. Further, at high traffic levels, the high level of collisions leads to very inefficient usage of the communication channel, just at a time when high capacity is desirable. Complicated algorithms are sometimes employed to prevent the delays becoming excessive, especially in the case of low priority nodes.

In Japanese patent application JA 61 102 842 (Ricoh Co Ltd), a CSMA/CD network is proposed which uses frequency division multiplexing to provide a fixed control channel and a number of data transfer channels, each channel occupying a different frequency band. Two nodes wishing to set up a one-to-one communication path negotiate over the control channel to find a free data transfer channel and agree to communicate over this channel. In this way, a number of one-to-one communications can take place simultaneously, giving high priority nodes the required rapid access to the network (assuming that these nodes maintain a data transfer channel allocated to them at all times). However, it still takes an indefinite time to set up the one-to-one communication links. This is because, as with all CSMA/CD type communication channels, the time taken to communicate over the CSMA/CD control channel cannot be guaranteed and consequently negotiating the use of a data channel over the control channel takes an indefinite time. As a result, at least some nodes, say high priority nodes which make infrequent use of communications so that it would be uneconomic to allocate them a data channel permanently, could experience unpredictable delays before being allocated a channel. An example of such a node is an error monitoring node, designed to signal to the system when a particular parameter exceeds its permitted range (eg speed of a conveyor, temperature of a furnace, etc). This would not be acceptable in many network applications such as, for example, an industrial LAN.

An object of the invention is to provide a communication network which mitigates the problems of prior art networks.

Accordingly, the invention provides a communication network for communication between a plurality of nodes over a multi-channel communication medium, the network comprising a plurality of allocatable channels for data communication between nodes and a control channel for allocating the allocatable channels to individual nodes, characterised in that at least one of the channels operates under a token-passing protocol with a limited message packet length whereby selected messages may be sent between nodes with a predetermined maximum delay.

By constraining the message packet length on the token-passing channel, the token retention time at each node and, consequently, the token cycle time for that channel can be kept low. It is possible to have a limited message packet length on the token-passing channel while still allowing for the efficient transmission of longer messages because of the availability of the allocatable channels for the bulk of the transmission of information.

Thus, a network in accordance with the invention can provide a guaranteed maximum delay time for any node wishing to communicate selected, albeit short, messages over the communication medium. This guaranteed maximum delay can be kept low because the token-retention time at each node can be set to a low level on the token-passing channel as only short messages are sent on this channel. With prior art, single channel, token-passing networks, the token retention time must be kept large to accommodate nodes requiring to transmit a large amount of data. Splitting the data into smaller packets, one to be transmitted on each receipt of the token, is undesirable as this leads to very slow transmission of this data and complications to the protocol. Prior art multi-channel networks based on CSMA/CD protocols cannot provide a guaranteed maximum delay because of the nature of those protocols.

Preferably, a channel operating under a token-passing protocol with a limited message packet length is the control channel, and the message packets transmitted by the nodes over the control channel comprise messages for negotiating the allocation of the allocatable channels to the individual nodes. In this way a maximum guaranteed delay can be provided for the allocation of an allocatable channel for the transmission of data messages between nodes.

The message packets transmitted by the nodes over the control channel preferably also comprise data messages unconnected with negotiating the allocation of the allocatable channels, whereby urgent data messages can be communicated between nodes over the control channel without needing to wait for the allocation of an allocatable channel. By permitting the control channel to also carry information packets which are not concerned with allocation of the allocatable channels, a node which does not currently have a channel allocated to it can nevertheless transmit an urgent message (eg. an error message) without encountering the delay of setting up an allocated channel and without the need to set aside a channel specifically for urgent communications.

A dedicated channel, other than the control channel, which operates under a token-passing protocol with a limited message packet length may, however, be provided and message packets comprising data messages unconnected with negotiating the allocation of the allocatable channels may be transmitted by the nodes over this dedicated channel. In this way urgent data messages (eg. error messages) can be communicated between nodes over this dedicated channel without needing to wait for the allocation of an allocatable channel.

In the particular example of a communication network described later, more than two nodes are permitted to communicate on an allocatable channel, thereby establishing a sub-network on this channel. These nodes can communicate using any desired protocol over this channel, such as MAP or Ethernet, or a token-passing protocol with a limited message packet length as discussed above, in order to meet particular requirements of the nodes on this sub-network. This arrangement allows for a plurality of sub-networks to be established on a single physical communication medium in an easy and flexible manner. Thus a plurality of sub-networks may each operate on a separate allocatable channel, each sub-network operating under the same or a different protocol, or no protocol. There could, for example, be two or more channels in the network which operate under a token-passing protocol with a limited message packet length. One could be the control channel, and one could be a dedicated channel as mentioned above dedicated to serving high priority nodes which make infrequent use of communications.

If each channel operates in a respective frequency band, channel separation is possible in a manner which permits compatibility with existing single channel equipment to be maintained for a token-passing channel while providing the additional channels. The number of channels which can be provided will depend on the bandwidth of the communication medium. For compatibility and reliability reasons, the frequency band of the token-passing channel should preferably be fixed. However this need not be the case and the bandwidth and the frequencies of the channels could be variable so that they may be allocated in a flexible manner depending on the needs of the nodes on the network.

Preferably, one or more nodes each comprises a switchable communication adapter which is switchable between frequency bands, whereby said one or more nodes can communicate on a plurality of channels. An alternative would be to provide said one or more nodes with a plurality of fixed communications adapters, one for each channel. However, this would be a more expensive solution and would limit flexibility.

A node can additionally comprise a second communication adapter which is operable in at least the frequency band of a channel operating under a token-passing protocol with a limited message packet length, whereby said at least one node can communicate simultaneously on this token-passing channel and on another channel. Such nodes will operate efficiently since they are permanently communicating on the token-passing channel. The provision of a second communication adapter avoids potential delays and other complications which might result from a node leaving and re-entering the token-passing ring. If required, a node could comprise more than two channel adapters, for example if there are two or more token-passing channels and a plurality of allocatable channels as mentioned above.

If a node comprises only a single switchable communications adapter, it preferably maintains a ring maintenance flag for indicating whether ring maintenance has been performed during the last complete rotation of the token on the token-passing channel, a logical ring being formed by the passing of the token on the token-passing channel, and ring maintenance is performed if the ring maintenance flag is not set when the token is received by that node.

Not all of the nodes in the network need to operate on the allocatable channels. For example, a node which merely needs to issue very short messages (eg. indicating an error condition) could be permitted to do so by sending an information packet on a token-passing channel as mentioned above. Such nodes have the advantage that they are simple (and hence low cost) since they only require a single non-switchable communication adapter. Accordingly therefore, some nodes comprise a fixed frequency band communication adapter which operates solely in the frequency band of a channel operating under a token-passing protocol with a limited message packet length, whereby communication by means of the fixed frequency band communication adapter is solely on this token-passing channel.

The invention finds particular but not exclusive application in local area networks in which the communication medium is a broadband cable, and the communication adapters are modems. The invention equally finds application in other types of network using a multi-channel communication medium, such as fibre-optics based communication networks.

A particular example of a communication network in accordance with the invention is in the following with reference to the accompanying drawings in which:

- Figure 1 is a schematic diagram of a typical arrangement of nodes on a token-bus communication network;

- Figure 2 is a graphical overview of the architecture of an example of a communication network in accordance with the invention;

- Figure 3 is a schematic block diagram of the architecture of high performance network node for simultaneous transmission on two channels;

- Figure 4 is a schematic block diagram of the architecture of a network node having a single, switchable communication adapter;

- Figure 5 is a schematic block diagram of the architecture of low-cost network node, for transmitting on a single, fixed channel; and - Figure 6 is a schematic illustration of an example of channel allocation on a communication network in accordance with the invention.

Figure 1 is a schematic diagram of a typical arrangement of nodes on a token-bus communication network, such as is known in the prior art. For reasons of convenience of illustration, only five nodes are shown. In practice, the number of nodes may be substantially more, or may be less than five. The communication network shown has nodes 1 to 5 communicating via a communication medium (eg. a broadband cable) 6. Each node in the network may only transmit information to other nodes when it is in possession of a "token", sent to it by its predecessor node in the network. The predecessor of node number 2, for example, is node number 3. The token is a unique identifier which the nodes recognise as giving them authority to transmit.

When a node (say node 2) is in possession of the token, it may pass as many messages as it requires, subject to restraints imposed on it by the protocol under which the network is operating to the other nodes in the network, and it then passes the token on to its successor node (ie node number 1 is the successor of node number 2) by sending a message including the token to that node. This node (eg. node 1) then, similarly, transmits as required and then passes the token on to its successor node (the successor of node 1 is the node 5).

It can be seen, therefore, that the predecessor and successor relationships are arranged to define a logical ring such that the token is passed repeatedly around this logical ring. It can also be seen that the predecessor and successor relationship is primarily a logical, as opposed to a physical one. The communication medium here is assumed to be a single broadband cable 6 connected at one end, the so called head end 7, to a head end remodulator 8. The broadband cable normally used for token-bus LANs, with the repeaters and other equipment has a bandwidth of about 400MHz. In a single cable system such as that illustrated, the total bandwidth is split into high and low regions for transmission in opposite directions, and signals are transferred from one to another by the head-end remodulator 8. In a dual cable system, the full bandwidth is available in both directions and the head-end remodulator is replaced by a repeater. In a network with a fibre optics communication medium, the available bandwidth can be orders of magnitude larger. It is possible in a fibre optics network for the network to be arranged in a star formation as will be well known to a person skilled in the art.

Consider a LAN based on a token-passing protocol and having 100 nodes spread over a radius of 5 Km. Bearing in mind the geographical spread, the worst case propagation time from the head end to the most remote node is of the order of 35 microseconds. The maximum time taken to transmit a token from one node to the next is given by the worst-case propagation time plus the time allowed for the recipient node to respond to receipt of the token. This total is called the transmission time (TT). Another parameter of the network which directly affects performance is the maximum retention time (MRT) for which a node is permitted to retain the token before passing it on. Once these times are known, the maximum access time for the network is given by the following relationship:

Maximum access time = number of nodes x [TT + MRT].

An estimate of the ultimate performance which could theoretically be obtained from a network may be found by assuming that the operation of the nodes is arbitrarily fast by the use of sufficiently high processing speeds and large bandwidth. This leaves only the propagation time as the limiting factor; so for the network of 100 nodes and 5Km radius, the transmission time may be about 35 microseconds and the token retention time may be taken as zero, giving a worst case maximum access time of 3.5 milliseconds.

In practice, real networks do not approach this limit, for two reasons. Firstly, bandwidths are not arbitrarily wide. The standard channel bandwidth in a broadband system is 12 MHz, giving a data rate of 10Mbit/second. The transmission of 1K bytes, whether data or protocol overhead, requires 819.2 micro-seconds. In order to allow the transmission of large files of data without excessive fragmentation, the maximum token retention time must be set to a large value, typically 5 to 10 milliseconds. Secondly, with the node operation taking a finite time, the transmission time comes out at about 40 microseconds. With these values, the 100 node, 5 KM network falls short of the theoretical ultimate performance by more than two orders of magnitude, with the result that a MAP network of this size is unusable as a real-time control system such as an industrial LAN.

An industrial LAN having 100 nodes within a radius of 5 Km is described in the following as an example of a communication network in accordance with the invention, although it should be appreciated that the present invention is neither limited to industrial LANs, nor to networks of this size.

With a cable bandwidth of 400MHz, a single broadband cable could accommodate about 12, and a dual cable about 30 simultaneous channels, each carrying about 10Mbits per second. Once again, the figures for fibre-optics would be orders of magnitude greater. In prior art token-bus networks, this available bandwidth is not fully used. The invention capitalises on this available bandwidth by providing a plurality of allocatable channels on which message communication between nodes can take place, in addition to a control channel. As the bulk of the data to be transferred between nodes can be sent over the allocatable channels, the individual message length on the control channel can be limited to a low value and the token retention time for the nodes, and consequently the total token retention time can be kept short, increasing the rate of response of the communication network.

The characteristics of this example of an industrial LAN are summarised in Table 1 below.

TABLE 1

Maximum propagation time (as before)	35 us
Transmission time	40 us
Control channel max message length	75 bytes
Control channel max token retention time	60 us
Token transfer time per node	100 us
Guaranteed access time for 100 nodes	10 ms max

With a 10 millisecond guaranteed access time, it is possible for each node on the network to transmit a short message 100 times per second, which is sufficient for the control of industrial machinery.

Figure 2 is a graphical description of an architecture for this example of a communication network in accordance with the present invention. Here, the structure of the Open Systems Interconnection (OSI) model has been adopted for both the token-passing channel and the allocatable channels. This Figure should be self explanatory to anyone versed in the OSI standards, and accordingly it will not be discussed further herein. The token-passing channel adheres to OSI standards in order to retain compatibility with existing LAN equipment based on those standards. As, however, the allocatable channels are allocated dynamically by negotiation amongst the nodes, any protocol, or no protocol, may be used on these channels once allocated, provided that there is a mechanism for closing down the channel on completion, to make it available for re-allocation.

Figure 3 illustrates the architecture of a node 20 as incorporated in this example of a communication network in accordance with the invention. The node comprises the following components:

- a switchable channel adapter, or switchable modem 22, capable of operating over a range of frequencies, preferably the full range of frequencies provided by the communication medium, and a second channel adapter, or fixed modem 24 operable on the frequency of the control channel 26;
- logic, in a communication controller 28 for implementing the protocol chosen for the control channel 26 - in this case IEEE 802.4 - and for associating it with the second modem 24;
- logic for implementing whatever protocol or protocols are to be used by the station on the allocated channels 30, 30′, 30″,... and for associating them with the switchable modem 22;
- supervisory logic, also in the communication controller 28 which recognises the need for a channel and uses the control channel protocol to negotiate the provision of one;
- the task (machine 32) operating at the station and initiating or responding to transmissions on the network.

A station may, for example, have the duty of handling the downloading of part programs to a numerically controlled machine tool (ie. machine 32 comprises the machine tool and appropriate processing logic) from a central database (not shown), whilst periodically transmitting a reading of the cutting tool temperatures to a remote monitoring program. It may also be required to handle occasional time-critical messages such as notification of loss of coolant supply. For the first task, an appropriate protocol would be MAP FTAM, while for the second, a simple serial point-to-point message service would suffice. The communication controller 28 in this station is able, therefore, implement both of these protocols and is able to associate either of them with an allocatable channel 30, 30′, 30″,... by means of the switchable modem 22. It is also able to implement IEEE 802.4, to associate it with the control channel 26 by means of the second modem 24, and to use it to acquire channel(s) from time to time. For time-critical messages, the station is able to send short messages on the control channel 26 using the second modem 24. In Figure 3, and also Figures 4 and 5, a double connecting line is representative of paths/channels for data and control information and a single connecting line is representative of a control path. Thus, the internal paths 34, 36 and 38 are for the passing of data and control information between the functional units indicated in the station 20, and the control path 40 is used for channel selection by the switchable modem 22.

The communication controller 28 (also the communication controllers 28′ and 28″ in Figures 4 and 5 respectively) are typically implemented in the form of software logic running on general purpose computing hardware such as a microprocessor and associated memory, etc. Also, in each of those figures, the communication controller and the machine may either be separate from one another, or they may be separate logical entities in the same physical device. For example, they may be separate tasks operating in a general purpose computer.

With the architecture illustrated in Figure 3, it possible for the node 20 to remain a member of the token-passing ring even whilst it is using another channel to transfer data. Whereas the switchable modem is frequency agile in order that it may be switched from channel to channel, the second modem 24 is permanently allocated to the control channel and consequently need not be frequency-agile. High-priority short messages can be transmitted on the control channel without interrupting the transmission or reception of a low-priority long message on an allocated channel.

Figure 4 is a schematic illustration of a lower cost alternative to the Figure 3 architecture for a network node. The node 42 shown in Figure 4 comprises a single, switchable modem 22. It does not include the second modem shown in Figure 3. This modem 22 allows communication to be made on only one channel at any one time, but is capable of switching between channels. The switchable modem in Figure 4 is, in the normal state, set to the frequency of the control channel 26 and consequently its node forms part of the token ring on that channel. On allocation of its node to an allocatable channel, the communication controller in the node 42 causes the frequency of the modem to be changed by means of control signals on path 44 to that of the allocated channel (ie. one of channels 30, 30′, 30″,...) and drops out of the token-passing ring in the control channel. The provision of a single modem of this type forms a relatively low cost approach to attaching nodes to the network, since only a single modem is needed. The sequence of events involved in allocating a channel, which is described below, is the same for node architectures illustrated in both Figure 3 and Figure 5. However, in the case of a single switchable modem, some delay can be experienced in rejoining the token ring on giving up the allocated channel. This is due to characteristics of IEEE 802.4 as explained later.

It may also be desired to provide some nodes on a network in the form of a simple node 46 such as that shown in Figure 5, with a single non-switchable modem 48. For example, a node which merely needs to issue very short messages (eg. indicating an error condition) may not need to use an allocatable channel. Such nodes may be implemented by means of an architecture as shown in Figure 5. The communication controller 28″, having determined that an error message is to be sent, transmits the error message as a short message on a channel designated for this purpose. The channel for this purpose could be the control channel, or it could be a dedicated channel especially set aside for the error messages.

It will be appreciated that in this way it is possible to retain full plug compatibility with single channel equipment such as that shown in Figure 5, from a prior art token-protocol network based on the Manufacturing Automation Protocol (MAP).

The simplest procedure is to allow the prior art equipment to communicate solely on the control channel. The prior nodes would use the control channel for all purposes, including communicating with the new nodes. The presence of such prior equipment on the network will, however, reduce the performance of the network as the token retention time will have to take account of the amount of data which a node with such equipment will put on the token-passing channel.

Another approach is to keep the new control channel separate from the old operating channel, and use the protocol to allocate a channel to the old equipment. A gateway would then have to be provided to interconnect the two sets of nodes; the gateway computer being responsible for requesting and maintaining the allocated channel. This would allow the new equipment to realise the full potential performance of the protocol, while the old would operate as before on a logically separate LAN on the same physical communication medium.

It will be appreciated that it is possible for a node to comprise more that two channel adapters. An example of a case where it may be desirable for there to be a third channel adapter in addition to the two shown in Figure 3, is where a channel separate from the control channel is specifically dedicated for error messages. This separate, dedicated channel could operate under a token-passing protocol as discussed above. Another example of a node where it may be desirable to have more than two channel adapters is in a node which serves as a gateway between different logical networks set up on the single physical communication medium.

The communication controller 28, 28′, 28″ in each station participating in the token-passing scheme on the control channel maintains information about the state of the network in the form of predecessor and successor node addresses and a channel status table.

The predecessor and successor node addresses are maintained in accordance with IEEE 802.4. Where a node has two modems, as illustrated in Figure 3, one for connection solely to the control channel and one for the allocatable channels, or a single non-switchable modem permanently allocated to the control channel as illustrated in Figure 5, these addresses are not used for any other purpose than ring maintenance as in IEEE 802.4. Where, however, there is a single switchable modem as illustrated in Figure 4, the addresses are also used for initiating action to take account of changes when a neighbouring node leaves the channel.

The channel status table (not shown) is a compact table which is preferably implemented in memory associated with the communications controller. The channel status table only requires one bit per channel for storing the occupancy condition of each of the channels which may be allocated. It is maintained by communications controller, which interprets all the channel allocation messages passing over the network, regardless of their destination address.

The protocol used on the control channel in this example is as specified by IEEE standard 802.4, with the additional requirement that the maximum length of any transmission, including token transfer, on the control channel is limited to a low value (eg. the 75 bytes referred to in Table 1). This restriction is made possible because the transmission of longer messages is removed from the control channel to one of the allocated channels; and it does not impair compatibility with other OSI equipment because, if a node on the network exceeds the limit value, it will not cause an error, it will simply reduce the extent of performance improvement.

The message packets or frames transmitted on the control channel in this example have the fields indicated in Table 2, below.

TABLE 2

Field No	Length (bytes)	Field Name
1	3	Preamble
2	1	Start Delimiter (SD)
3	1	Frame Control (FC)
4	6	Destination Address (DA)
5	6	Source Address (SA)
6	34 max	Control Channel Data Unit (CCDU)
7	4	Frame Check Sequence (FCS)
8	1	End Delimiter (ED)
9	1	SD
10	1	FC = token identifier
11	6	Token Successor Address
12	6	SA
13	4	FCS
14	1	ED
Total	75 max	

It should be noted that 6 byte addresses are used to comply with the OSI global addressing standard. It will be noted also that the message packet shown in Table 1 is effectively two messages concatenated together; hence the two start delimiters and the two end delimiters. Fields 9 to 14 are omitted in those cases where the token is not passed on immediately after the message is transmitted. In general this will only occur during ring maintenance.

The Control Channel Data Unit (CCDU) contains the fields indicated in Table 3 below.

TABLE 3

Field No	Length (bytes)	Field Name
	EITHER	
1	1	Data unit Type
2	1	Channel Number (CN)
3	6	Token predecessor address
4	26 max	Allocated channel protocol data
	OR	
1	1	Data Unit type
2	33 max	Unrestricted data
Total	34 max	

The Data Unit Type field determines whether the Data Unit is a Channel Allocation Request, a Channel Allocation Response, a Chemical Release Request, a Channel Release Response, a Channel Abort or a Short Message. In the first five cases, a second field contains the channel number concerned, and the third field contains the address of the source node's token predecessor, for use in passive ring maintenance, if needed.

This packet format is entirely within the constraints of IEEE 802.4, with the additional constraint of a very low length limit.

The functions carried out by the participants in the control channel include those required by IEEE 802.4 and additionally those needed to maintain the channel status table.

The communication controller of each participant in the control channel not currently in possession of the token listens to the traffic on the bus for maintaining the internal tables of information about the current state of the network.

The communication controller maintains a record of the addresses of its predecessor and successor in the ring as required by IEEE 802.4. Whenever one of these addresses appears in a Channel Allocation Request or Response, it is assumed that that node is leaving the ring. Accordingly, the address is replaced by its predecessor or successor address, obtained from the Channel Allocation Request or Response message.

The communication controller node additionally maintains its channel status table showing the occupancy of each of the allocatable channels. Whenever a channel number appears in a Channel Allocation Request that channel is flagged as occupied, and when a channel number appears in a Channel Release Response or a negative Channel Allocation Response message, that channel is flagged as unoccupied, respectively in the appropriate location in the channel status table. These terms will be explained later.

In the case of a node having a single, switchable modem and no second modem, when this node leaves the ring to use another channel, the table becomes inaccurate; so on rejoining the ring, the table is restored using information transmitted by the admitting note. This may be performed using the ring maintenance procedures of IEEE 802.4. However, as mentioned above, some delay can be experienced at this stage.

Channel negotiation/short message functions are carried out whenever the node wishes to use the network to carry data, either by the allocation of a channel, or by using the control channel to carry a message in the CCDU field.

A short message is a single packet on the control channel, which may or may not require a response. It contains the fields indicated in Table 4, below.

TABLE 4

Field No	Length (bytes)	Field Name
1	3	Preamble
2	1	Start Delimiter (SD)
3	1	Frame Control (FC)
4	6	Destination Address (DA)
5	6	Source Address (SA)
6	1	Data Unit type = Short Message
7	33 max	Message text
8	4	Frame Check Sequence (FCS)
9	1	End Delimiter (ED)
10	1	SD
11	1	FC = token identifier
12	6	Token Successor Address
13	6	SA
14	4	FCS
15	1	ED
Total	75 max	

When a node wishes to use an allocatable channel for a point-to-point link the following sequence of operations is performed.

1. The initiating node consults its channel status table to identify an unoccupied channel. If none exists, it waits until one becomes available.

2. On receiving the token, the node issues a Channel Allocation Request message, addressed to the intended recipient and naming the channel, and immediately passes on the token. Other nodes will then interpret this packet and identify that channel as occupied. There is no possibility of the same channel being requested simultaneously by two nodes, since only the node in possession of the token can transmit. If there were no unoccupied channel available, the initiating node could, on receiving the token, send a Short Message to the recipient to indicate that a long message is pending.

3. The recipient, after receiving a Channel Allocation Request message, and recognising its address in that message, waits until it receives the token, and then transmits a Channel Allocation Response message, which may be positive or negative to indicate acceptance or rejection, to the initiator. If the response is positive, the recipient switches its switchable modem to the allocated channel. This response message will be interpreted by other nodes, and used either to confirm or cancel the occupied status of the channel. The initiator, on receiving the response message, switches its switchable modem to the named channel and begins transmission.

In the case of a single modem, if for any reason communication is not established on the allocated channel within a time-out period, and in any case after completion of the communication on the allocated channel, both initiator and recipient return to the control channel and rejoin the ring.

The following sequence of messages is issued to release a channel when a point-to-point link is no longer needed.

1. Either of the participants in the link transmits a Channel Release Request message addressed to the other.

2. On receipt of that, and after transmitting any waiting traffic on the allocated channel, the other participant transmits a Channel Release Response message to the first participant. Other nodes will interpret this packet, and flag the named channel as unoccupied.

3. Additionally, as confirmation of the release of the channel, either participant may broadcasts a Channel Release message on the control channel, naming the channel just released. This confirmation is needed for the single modem case where the Channel Release transmissions in steps 1 and 2 are on the allocated channel.

As an alternative to the above procedure, either participant may transmit a Channel Abort message to the other. Both participants must then cease using the allocated channel immediately and discard any waiting traffic. The Channel Abort message indicates to other nodes that the channel is freed, and the optional confirmation procedure may be used. Once again, the confirmation is need for the single modem case to indicate over the control channel to the other nodes that the channel is freed.

When an allocatable channel is required for a multi-node network, using a protocol such as token-passing or CSMA/CD protocol, a different sequence of steps is performed. As has always been the case, every such network has a special station, usually called the Network Manager, which is responsible for administration of the network, directory services etc. The sequence of steps is as follows.

1. The node containing the Network Manager of a logical network, having found an unoccupied channel, transmits a Channel Allocation Request packet, followed by a Channel Allocation Response packet, both addressed to itself. This has the effect of reserving the channel and notifying other nodes that it is occupied.

2. Every node which subsequently wishes to take part in the logical network transmits a Channel Number Request packet addressed to the Network Manager.

3. The Network Manager responds with a Channel Number Response packet containing the number of the channel allocated for that logical network, and a code identifying the protocol to be used on that channel.

4. The node then joins the network on the allocated channel, using whatever joining procedure is laid down by the protocol. It may subsequently leave the logical network in the same way.

The Channel Number Request packet is really a special case of the Channel Allocation Request packet, and may be realised as a Channel Allocation Request packet with an empty address field. Similarly, the Channel Number Response packet may be the same as the Channel Allocation Response packet.

On closing a logical network, once all the participants in a logical network except the Network Manager have left it, the Network Manager may, if appropriate, release the channel by transmitting Channel Release Request and Response frames addressed to itself; by transmitting a Channel Abort packet to itself; or simply by broadcasting a Channel Release packet.

Figure 8 is a schematic illustration of an example of channel allocation on a communication network in accordance with the invention. It is assumed in this example that there are 100 nodes on the network as described above, the nodes being numbered 0 to 99. It will be noted that in addition to the control channel, there are channels currently unallocated and channels supporting communications under a MAP protocol, an Ethernet protocol and communications on a point to point basis. The allocation of these channels was made as described above (ie. under the token-passing protocol on the control channel). It should be noted that Figure 8 is purely schematic, and the width of the channels shown does not represent their relative bandwidths. The bandwidth of each of the channels could be the same. Equally, channels of differing bandwidths may be provided and the channels may be allocated in accordance with bandwidth requirements for the communications to be performed. This may be done by preselecting a plurality of channels of different bandwidths, or by dynamically determining the bandwidths of the channels.

A particular example of a communication network in accordance with the invention has been described above, in the form of an industrial Local Area Network (LAN). However, it will be apparent that the invention is applicable to other environments, such as satellite transmission, or wide area networks. Accordingly, the invention is not limited to LANs, and particularly not limited to the LAN described herein. The communication medium of the preferred example described above is a broadband cable. However, it will be appreciated that the invention is applicable to communication networks using other communication media, such as electromagnetic communication through free space, fibre optic cables and so on, where the communication medium be divided into channels (eg. by means of frequency division multiplexing (FDM)), and all messages are broadcast into the transmission medium in such a way that they can be received by all receivers connected to the medium.

It will be appreciated that further modifications and additions to the communication network described in the above are possible within the scope of the present invention. For example, in the particular example described above, dual modems are used, the control channel protocol is IEEE 802.4, unacknowledged and the message length limit is set to 75 bytes. However, in another example of a communication network in accordance with the invention:

- single switchable modem nodes may be used instead of, or in addition to the dual modem nodes of the above network; and/or
- a token-passing protocol not in accordance with IEEE 802.4 may be used on the control channel; and/or
- the message length limit in the control channel may be set to another value, chosen to optimise performance in particular circumstances; and/or
- acknowledged, rather than unacknowledged messaging may be used in the control channel.

In summary, the invention provides a communication network which:

- enables wide compatibility to be achieved because the control channel can be arranged to adhere to an existing protocol (for example, in the network described above, software compatibility need not be

compromised because the high layers of the architecture adhere to the standards of the OSI model);
- enables high performance to be achieved because the bulk transfer of data is handled by allocated channels, and does not therefore interfere with the process of gaining access to the network and sending high priority short messages;
- enables a high degree of flexibility because its performance is high enough to enable the elimination of physically separate sub-network structures for real-time control, and affords the opportunity to allow equipment using non-standard protocols to be integrated into the network and brought under the control of the token-passing operating protocol; and
- affords the opportunity to increase reliability and maintainability by introducing redundancy into the network, in the form of distributed processing making use of the high performance.

In addition, a communication system in accordance with the invention provides the potential for bringing the total management of a cable system - not just the channels allocated to the computer network - under the control of the protocol. This is a great advantage where the cable is required to carry non-computer signals such as video or voice communications. Instead of the permanent allocation of channels to such purposes, with consequent inability to remove them rapidly in case of equipment malfunction, the channels are allocated dynamically and can be managed by the nodes in the network and/or under the control of a controlling computer.

Some ways of making use of the facilities which can be provided in a communication network based on the invention are:
- point-to-point communications on a control channel, in which each data transmission requires the allocation of a channel between the sender and every intended recipient;
- multi-network configurations, in which multiple logical networks, all using the same protocol but each on a separate channel, are established on the common communication medium.
- multi-protocol configurations, in which multiple logical networks, using different protocols and on separate channels, are established on the common transmission medium.
- short message transmission on the control channel for:

(i) rapid transmission of high priority messages;

(ii) communicating with a node or nodes which are not part of a particular logical subnetwork in order to provide a measure of cross-network communication (Note that this provides none of the data transformation features of OSI, so it requires individual handling by the application program initiating the communication);

(iii) communication between network management agents and a network manager (most OSI protocols call for network management agents to be present in each of the nodes on a network); - the provision of some low cost nodes operating only on the control channel, for connecting simple sensors or actuators to the network in such a way that they could be addressed by a data collector/distributor which is simultaneously a member of, say, a logical MAP network;
- coping with errors on a logical subnetwork (eg. if a node fails in 'jabber' state - ie, constantly transmitting noise on to the network - the fault can be located and cured by moving nodes one by one to another channel).

In the above specific example of a communication network according to the invention, ring maintenance is performed in accordance with IEEE 802.4. This can cause delays when nodes having a single switchable modem wish to rejoin the token-passing channel on releasing an allocated channel. This is because the updating of the predecessor and successor information for adding a node to the token-passing channel takes place during ring maintenance, and this does not occur during every token ring cycle under IEEE 802.4. Ring maintenance is optimised in accordance with IEEE 802.4 on the assumption that long messages as well as short ones are being transmitted on the token-passing channel. In the case of a token-passing channel having messages with a limited message package length in a network in accordance with the invention, yet further improved performance can be obtained for nodes having a single switchable modem by adopting a revised ring maintenance strategy. By adopting the strategy outlined below, the time taken between ring maintenance operations can be reduced.

As a part of this revised strategy, each node maintains a flag called the Ring Maintenance flag, or RM flag, which indicates whether ring maintenance has been performed during the last complete rotation of the token. The RM flag is reset on relinquishing the token, except when ring maintenance has been carried out by the relinquishing node, and is set when ring maintenance is observed to be carried out by another node.

Ring maintenance under this revised strategy is carried out :
- if no response is heard within one slot time after passing on the token (an error condition); or
- if the RM flag is not set when the token is received.

This ensures that, in the absence of errors, it is carried out by each node in turn, with a complete token rotation between occurrences. The principal aim of this is to make the time taken to rejoin the ring more predictable.

The functions carried out in accordance with this ring maintenance strategy are substantially the same as those of the 'solicit successor' procedure of IEEE 802.4, which will be known to one skilled in the art. One extra function is carried out in addition to the solicit successor procedure, however. That is the transfer of the Channel Status Table, whenever a node has been admitted to the ring, to the newly admitted node. The table can be sent as a Short Message to the new node, at the same time that the token is passed on.

Claims

1. A communication network for communication between a plurality of nodes (1-5) over a multi-channel communication medium (6), the network comprising a plurality of allocatable channels (30,30′,30″) for data communication between nodes (1-5) and a control channel (26) for allocating the allocatable channels (30,30′,30″) to individual nodes (1-5), characterised in that at least one of the channels (26,30,20′,30″) operates under a token-passing protocol with a limited message packet length whereby selected messages may be sent between nodes (1-5) with a predetermined maximum delay.

2. A network as claimed in claim 1, wherein a channel operating under a token-passing protocol with a limited message packet length is the control channel (26), and wherein the message packets transmitted by the nodes (1-5) over the control channel (26) comprise messages for negotiating the allocation of the allocatable channels (30,30′,30″) to the individual nodes (1-5).

3. A network as claimed in claim 2 wherein the message packets transmitted by the nodes (1-5) over the control channel (26) also comprise data messages unconnected with negotiating the allocation of the allocatable channels (30,30′,30″), whereby urgent data messages can be communicated between nodes (1-5) over the control channel (26) without needing to wait for the allocation of an allocatable channel (30,30′30″).

4. A network as claimed in any one of the preceding claims wherein a dedicated channel other than the control channel (26) operates under a token-passing protocol with a limited message packet length and message packets comprising data messages unconnected with negotiating the allocation of the allocatable channels are transmitted by the nodes (1-5) over this dedicated channel, whereby urgent data messages can be communicated between nodes (1-5) over this token-passing channel without needing to wait for the allocation of an allocatable channel (30,30′30″).

5. A network as claimed in any one of the preceding claims wherein more than two nodes (1-5) communicate on an allocatable channel (30,30′30″), thereby establishing a sub-network on this allocatable channel (30,30′30″).

6. A network as claimed in claim 5 in which the sub-network operates under a suitable network protocol.

7. A network as claimed in claim 6 wherein a plurality of sub-networks each operate on a separate allocatable channel (30,30′30″) each sub-network operating under the same or a different protocol, or no protocol.

8. A network as claimed in any one of the preceding claims wherein each channel (26,30,30′30″) operates in a respective frequency band.

9. A network as claimed in claim 8 wherein the bandwidth one or more of the channels (26,30,30′30″) is selectable.

10. A network as claimed in claim 8 or claim 9 wherein one or more nodes (1-5,20;42) each comprises a switchable communication adapter (22) which is switchable between frequency bands, whereby said one or more nodes (1-5,20;42) can communicate on a selectable one of a plurality of channels (26,30,30′30″).

11. A network as claimed in claim 10 wherein at least one node 1-5;20) additionally comprises a second communication adapter (24) which is operable in at least the frequency band of a channel (26) operating under a token-passing protocol with a limited message packet length, whereby said at least one node (1-5;20) can communicate simultaneously on this token-passing channel (26) and on another channel (30,30′,30″).

12. A network as in any one of claims 8 to 11 wherein some nodes (1-5;20) comprise a fixed frequency band communication adapter (46) which operates solely in the frequency band of a channel (26) operating under a token-passing protocol with a limited message packet length, whereby communication by means of the fixed frequency band communication adapter (46), is solely on this token-passing channel (26).

13. A network as claimed in any of the preceding claims wherein the communication adapters (22,24,46) are modems.

14. A network as claimed in claim 10 wherein a node (1-5;42) maintains a ring maintenance flag for

indicating whether ring maintenance has been performed during the last complete rotation of the token on the token-passing channel (26), a logical ring being formed by the passing of the token on the token-passing channel and wherein ring maintenance is performed if the ring maintenance flag is not set when the token is received by that node (1-5:42).

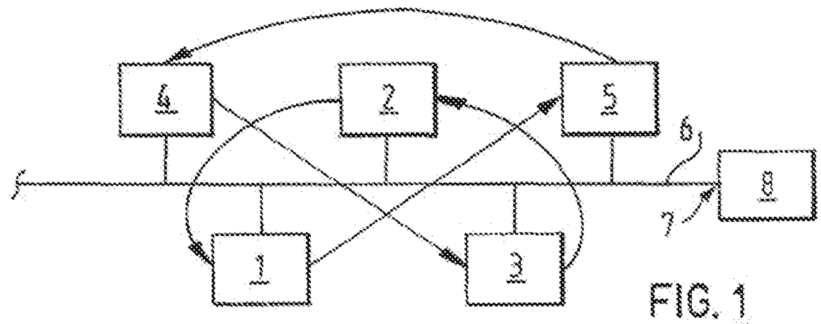

FIG. 1

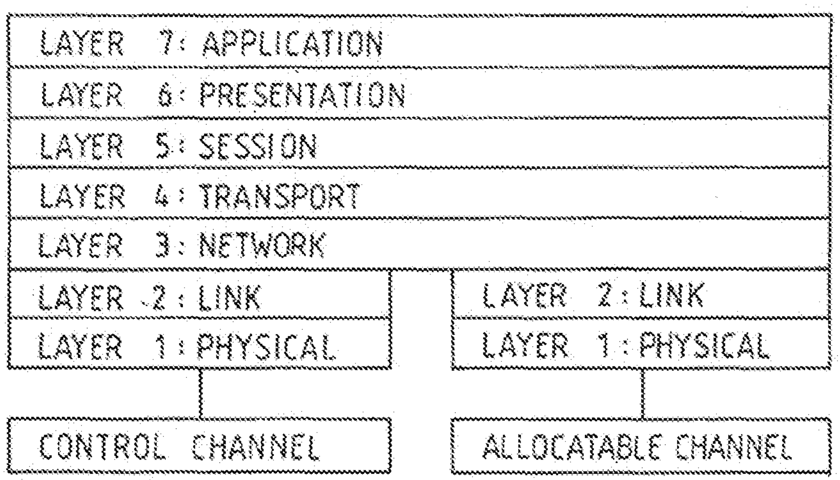

FIG. 2

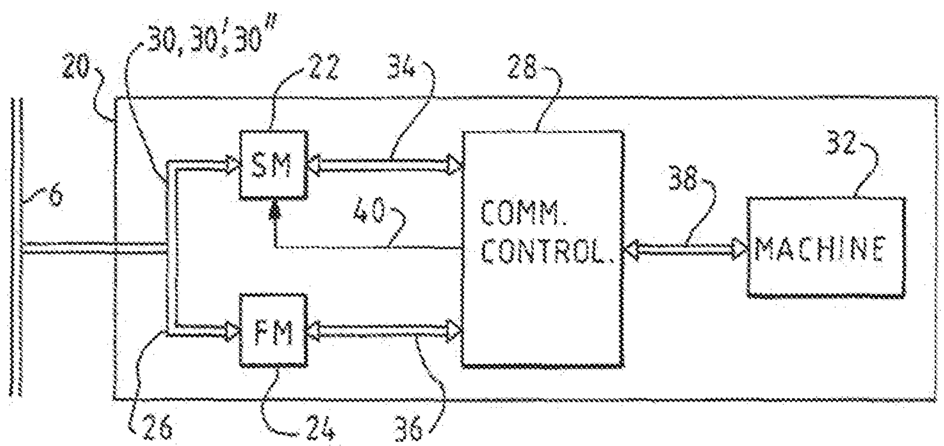

FIG. 3

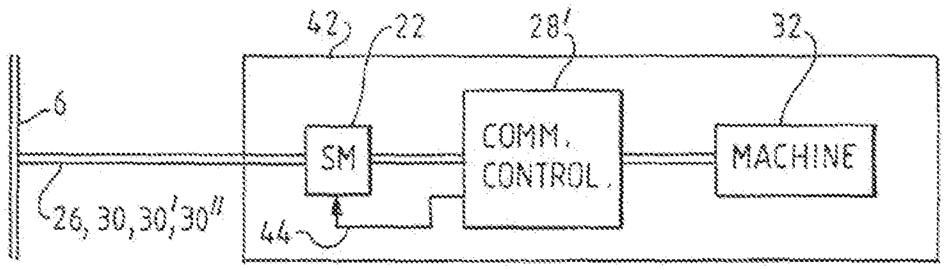

FIG. 4

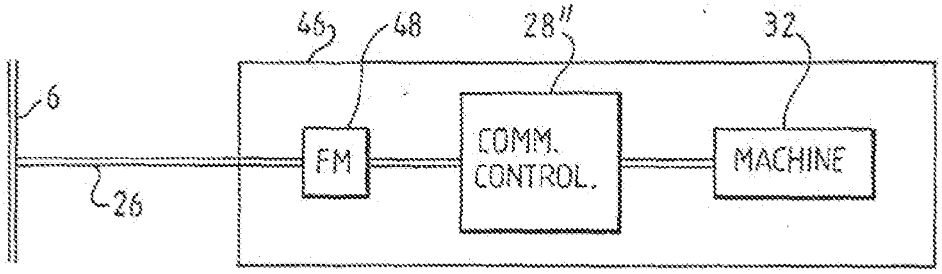

FIG. 5

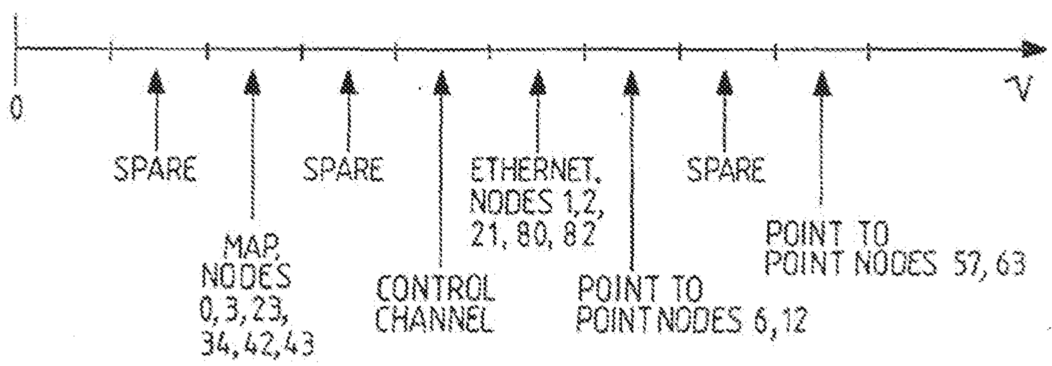

FIG. 6

European Patent Office

EUROPEAN SEARCH REPORT

Application Number

EP 88 30 9854

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cl.3)
A	THE 6TH INTERNATIONAL CONFERENCE ON DISTRIBUTED COMPUTING SYSTEMS, 19th-23rd May 1986, Cambridge, Massachusetts, pages 206-215, IEEE, New York, US; S. THANAWASTIEN et al.: "A CSMA/CD-token ring hybrid architecture for local area networks" * Page 206, left-hand column, lines 6-20; page 206, right-hand column, lines 10-12; page 210, right-hand column, lines 22-34 *	1-4,11	H 04 L 12/28
A	38TH IEEE VEHICULAR TECHNOLOGY CONFERENCE, 15th-17th June 1988, Philadelphia, Pennsylvania, pages 44-51, IEEE, New York, US; A.M. GLASS et al.: "Broad-band networks for lan and radio applications" * Page 45, left-hand column, lines 40-44; page 46, left-hand column, lines 24-32 *	1,2,5-10	TECHNICAL FIELDS SEARCHED (Int. Cl.3)
A	IEEE INTERNATIONAL CONFERENCE ON COMMUNICATIONS '86, 22nd-25th June 1986, Toronto, vol. 2, pages 931-933, IEEE, New York, US; M.S. GOODMAN et al.: "Application of wavelength division multiplexing to communication network architectures" * Page 931, right-hand column, lines 5-34 *	1,2,5-10	H 04 J H 04 L

The present search report has been drawn up for all claims

Place of search	Date of completion of the search	Examiner
THE HAGUE	09-06-1989	DE LA FUENTE DEL AGUA R.

CATEGORY OF CITED DOCUMENTS

X : particularly relevant if taken alone
Y : particularly relevant if combined with another document of the same category
A : technological background
O : non-written disclosure
P : intermediate document

T : theory or principle underlying the invention
E : earlier patent document, but published on, or after the filing date
D : document cited in the application
L : document cited for other reasons

& : member of the same patent family, corresponding document

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 866 590 A2

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
23.09.1998 Bulletin 1998/39

(51) Int. Cl.[6]: **H04L 29/06**, G06F 9/46

(21) Application number: 98104347.4

(22) Date of filing: 11.03.1998

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 13.03.1997 US 816914

(71) Applicant: MOTOROLA, INC.
Schaumburg, IL 60196 (US)

(72) Inventors:
- Whalen, Jon S.
Arlington Heights, Illinois 60005 (US)
- Whittington, David
Austin, Texas 78739 (US)
- Zabolotzky, Scott
Elgin, Illinois 60120 (US)
- Zurcher, Rodd
Lake Zurich, Illinois 60047 (US)
- Biersach, David
Lake in the Hills, Illinois 60102 (US)

(74) Representative:
Morgan, Marc et al
Motorola European Intellectual Property Operations,
Midpoint,
Alencon Link
Basingstoke, Hampshire RG21 7PL (GB)

(54) **System and method for delivery of information over narrow-band communications links**

(57) System and method for delivery of information over narrow-band communications links. The system has at least a browser (12), a mobile client (10), a fixed server (30) and an origin host (50). The browser (12) requests a resource. The mobile client (10) transmits the request to the fixed server (30). The fixed server (30) retrieves a primary resource from the origin host (50) and any dependent resources. The fixed server (30) transmits the primary resource to the mobile client (10). The mobile client (10) transmits an acknowledgment list to the fixed server (30) requesting certain dependent resources and sends the primary resource to the browser (12). The fixed server (30) transmits the requested dependent resources to the mobile client (10) in one transmission. The mobile client (10) sends the dependent resources to the browser (12) upon request. Thus, only transmitting two round-trips of data across the narrow-band communications link to transfer all the necessary data to create an entire information page reduces the delay significantly.

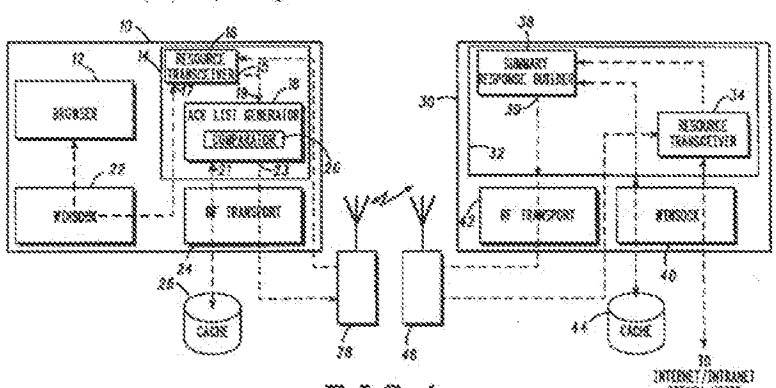

FIG. 1

EP 0 866 590 A2

Description

Field of the Invention

This invention relates to efficient delivery of information to browser clients over wide-area, narrow-band communications systems, including, but not limited to, Packet Data Networks and Circuit Switched Networks.

Background of the Invention

The World Wide Web (Web) is an ubiquitous communications network used to readily access available resources on many computers throughout the world and is attached to at least one computer network known as the Internet. The Web comprises a body of software, a set of protocols and a set of defined conventions for obtaining information on the Web. The Web utilizes hypertext and multimedia techniques to make the Web "user-friendly" for anyone who desires to browse, roam or contribute to the Web.

A HyperText Transport Protocol (HTTP) is a protocol that is used for transporting hypertext files across the Internet. In ordinary HTTP operation, a proxy receives an HTTP request for a resource and connects to a host identified in a uniform resource locator (URL). A URL is a standardized way of representing different documents, media and network services on the Web. The proxy retrieves the resource and returns a HTTP response to the requester.

In normal HTTP operation, the browser requests a HyperText Markup Language (HTML) response. A HTML is a standardized way to creak hypertext documents for use on the Web; HTML is a coding language that surrounds the text used in the hypertext documents with codes and brackets to indicate how the text should appear to the user. When the browser receives the HTML response, the browser parses it and issues individual requests for dependent resources, such as in-line images. Over a narrow-band, high latency connection, this "ping-ponging" (e.g., the browser requesting and receiving each dependent resource individually) results in a severe delay in completing the retrieval of an entire Web page. For a page containing N in-line resources retrieved over a link with an average round-trip latency of L seconds, the delay is approximately ((N+1)*L) seconds.

As a result of delays in delivery of information over narrow-band communications links, there exists a need for a system that delivers information to browser clients over wide-area, narrow-band communications systems in an efficient manner.

A preferred embodiment of the invention, is now described, by way of example only, with reference to the drawings.

Brief Description of the Drawings

The features of the present invention are set forth with particularity in the appended claims. Preferred embodiments of the invention are now described, by way of example only, with reference to the accompanying drawings in which:

FIG. 1 is an overview diagram of a wireless Web proxy system according to a preferred embodiment of the invention;
FIG. 2 is a bounce diagram of the preferred embodiment of the invention;
FIGS. 3 and 4 together are a flow chart of an operation of a mobile client according to the preferred embodiment of the invention;
FIGS. 5 and 6 together are a flow chart of an operation of a fixed server according to the preferred embodiment of the invention;
FIG. 7 is a flow chart of a method of operation for a cache check according to the preferred embodiment of the invention; and
FIG. 8 is a bounce diagram of an alternative embodiment of the invention.

It will be appreciated that for simplicity and clarity of illustration, elements shown in the figures have not necessarily been drawn to scale. Where considered appropriate, reference numerals have been repeated among the figures to indicate corresponding elements.

Detailed Description of the Drawings

A wireless Web proxy system, which is now described, is a system of middleware software which functions as a HTTP proxy incorporating a proprietary protocol for communication of HTTP requests and responses. The wireless Web proxy system provides the means for efficient delivery of information (e.g., Web content) such as text, images, sounds, and other resources, accessed via the HTTP protocol over Transport Control Protocol/ Internet Protocol (TCP/IP) networks (Intra-/ Internet) to browser clients over wide-area, narrow-band communications systems, including, but not limited to, Packet Data Networks (e.g., DataTAC 4000/ 5000/ 6000, Mobitex, CDPD, etc.) and Circuit Switched Networks (e.g., Analog Cellular, GSM, etc.).

As shown in FIG. 1, the wireless Web proxy system consists of mobile proxy software 14 installed on a digital processor/ mobile client 10 and fixed proxy software 32 installed on a digital processor/ fixed server 30. The mobile proxy software 14 implements an interface which conforms to the specification contained in RFC 1945, "Hypertext Transfer Protocol -- HTTP/1.0" for an HTTP 1.0 complaint proxy server. The digital processor 10 comprises at least the following components: a browser 12, the mobile proxy software 14, a winsock 22 and a radio frequency transport 24. The mobile proxy software

14 further comprises at least a resource transceiver 16 and an acknowledgment list generator 18 having at least a comparator 20. Connected to the mobile proxy software 14 of the digital processor 10 is a cache 26 and a radio transceiver 28. The wireless Web proxy system further comprises a memory having instructions and data stored therein that, when executed, causes the digital processor 10 and the cache 26 to comprise the resource transceiver 16 with an input 17 couplable to the browser 12 and an output 15. The acknowledgment list generator 18 has a first input 19 coupled to the output 15 of the resource transceiver 16, a second input 21 coupled to the cache and an output 23 couplable to the radio transceiver 28. Such a configuration allows the comparator 20 to compare the received resources with the cached resources.

The fixed proxy software 32 is installed and executed on the fixed server PC running Windows NT 3.51 or later. The fixed proxy software 32 implements the HTTP 1.0 client protocol and is responsible for retrieving resources from HTTP servers (Web servers) on the Internet or intranet. The digital processor 30 comprises at least the following components: the fixed proxy software 32, a winsock 40 and a radio frequency transport 42. The fixed proxy software 32 further comprises a resource transceiver 34 and a summary response builder 38. Connected to the fixed proxy software 32 of the digital processor 30 is a cache 44 and a radio transceiver 46. The wireless Web proxy system further comprises a memory having instructions and data stored therein that, when executed, causes the digital processor 30 to comprise the resource transceiver 34 and the summary response builder 38. The summary response builder 38 is coupled to the resource transceiver 34. The summary response builder 38 has an output 39 which provides a summary response which includes a plurality of status codes corresponding to a plurality of requested resources and includes the requested resources when available. The summary response comprises contents from a plurality of resources which when taken together with resources cached locally at the digital processor on the mobile client constitutes an entire information page (e.g., a Web page). In addition, the radio transceiver 46 is coupled to the output 39 of the summary response builder 38 for sending the summary response over a communications link to the browser 12.

A mobile user begins browsing by launching the mobile proxy software 14. This automatically launches the user's preferred Web browser software. The mobile user can browse the Web by entering uniform resource locators (URLs) by following links as he/she is normally accustomed to doing in a wire-line environment (e.g., on a Local Area Network).

FIG. 2 is a bounce diagram of the preferred embodiment of the invention. In FIG. 2, the browser 12 and the mobile proxy software 14 are the primary components which comprise the mobile client. Also shown in FIG. 2 are the fixed proxy software 32 and the origin host (e.g., Web site) 50. Requests and responses are exchanged between the mobile client and the fixed server over a narrow-band communications link (e.g., the transmissions between the mobile proxy software 14 and the fixed proxy software 32). As shown, the wireless Web proxy protocol only requires transmitting two round-trips of data across the narrow-band communications link to transfer all the necessary data which when taken together with the resources cached locally at the mobile client constitutes an entire information page (e.g., Web page). By reducing data transfer over the narrow-band communications link to a total of two-round-trips, the delay is reduced to 2*L seconds, where the average round-trip latency over a narrow-band communications link is L seconds.

FIGS. 3 and 4 together are a flow chart of an operation of a mobile client 10 according to the preferred embodiment of the invention. The mobile user's browser 12 is configured to treat the mobile software as a Web proxy server. In FIGS. 3 and 4, when the mobile user opens a URL, the browser 12 submits an HTTP request to the mobile client 10. The mobile proxy software 14 receives the HTTP request and examines its local cache (a database of URL-indexed information) at steps 102 and 104. The manner in which the mobile proxy software 14 examines its local cache 26 is described below in conjunction with FIG. 7. The mobile client 10 determines at step 106 whether it can respond to the browser 12 immediately or whether it must ask for the information from the fixed server over the narrow-band communication link. If the mobile client 10 already has certain resources cached (e.g., already has received a page before expiration), then the mobile proxy software 14 informs the fixed proxy software 32 not to send the resources that are cached at the mobile client 10 and the mobile proxy software 14 sends the HTTP response to the browser 12 at step 108. However, if the mobile client 10 decides that it must forward the request to the fixed server 30 because the resource is not cached or the resource has expired, the mobile proxy software 14 transmits the request to the fixed server in the form of a tokenized and compressed HTTP request at steps 110 and 112. The manner in which the HTTP request is tokenized is described below. Thus, from step 112, the operation of the mobile client pauses before proceeding to step 126. During this time, certain steps are performed at the fixed server, as shown in FIGS. 5 and 6.

In FIGS. 5 and 6, the fixed server 30, upon receiving the HTTP request at step 114, expands the tokenized and compressed HTTP request at step 116. After expansion of the HTTP request, the fixed proxy software 32 inspects its local cache 44 at step 118 to determine whether it may have any version of the requested primary resource stored locally. Again, the manner in which the fixed proxy software 32 inspects its local cache 44 is described below in conjunction with FIG. 7. If the fixed server 30 determines at step 120 that the primary resource is not cached locally or if there is a version of the primary resource cached but that version of the primary resource has expired, the fixed proxy software 32 connects to the origin host 50 (e.g., Web site) identified in the primary URL (e.g., Web, Gopher, File Transfer Protocol, etc.) or to another proxy, to retrieve the primary resource requested by the mobile client 10 at step 122. After the fixed proxy software 32 receives a valid version of the primary resource either from step 118 or from step 125, the fixed proxy software 32 determines whether the HTTP response having the primary resource is in the form of a HyperText Markup Language (HTML) at step 126.

If the HTTP response is not in a HTML form, the fixed proxy software 32 immediately sends the HTTP response to the mobile client 10 at step 128 in a tokenized and compressed form. The manner in which the HTTP response is tokenized is also described below. However, if the HTTP response is in a HTML form, the fixed proxy software 32 examines the resource identified by the primary URL to determine whether the mobile client 10 may need any other resources ("dependent resources", identified by absolute or relative URLs in tags in the HTML page) in order for the browser 12 to completely display the primary resource to the mobile user at step 130 (refer to FIG. 5). For example, an HTML page may contain images, Java applets, sounds, or other dependent resources which must be available for the browser 12 to properly display the page. If there are dependent resources identified, the fixed proxy software 32 examines its local cache 44, and if necessary, reconnects to the origin host 50 identified in the primary URL (e.g., for resources identified by relative URLs) or to other hosts (e.g., for resources identified by absolute URLs) and issues requests for those dependent resources (i.e., the fixed proxy software pre-fetches the dependent resources identified from the primary resource) and receives an updated resource from the origin host 50 or from whichever host the fixed proxy software 32 requested the resource. If the cached version of the primary resource of the mobile client 10 is up-to-date, the fixed proxy software 32 returns an indication that the mobile client 10 has the current version of the resource. Otherwise, the fixed proxy software 32 transmits the primary resource (in the form of a compressed HTTP response), along with information identifying the dependent resources upon which the primary resource depends, to the mobile client 10 at step 134.

When the mobile client 10 receives the HTTP response at step 136 of FIG. 3, the mobile proxy software 14 expands the HTTP response and updates its local cache 26 with the primary resource at step 138. If the mobile proxy software 14 determines at step 140 that the HTTP response the mobile proxy software 14 received at step 136 is not in a HTML form, the mobile proxy software 14 immediately sends the HTTP response to the browser 12 at step 108. However, if the mobile proxy software 14 determines at step 140 that the HTTP response it received at step 136 is in a HTML form, the mobile proxy software 14 identifies the dependent resources at step 142. Using the information received from the fixed proxy software 32 regarding the dependent resources, the mobile proxy software 14 examines its local cache 26 to determine whether it has any or all of them. Based upon this cache check (as described below in conjunction with the discussion of FIG. 7), the mobile proxy software 14 constructs a short acknowledgment list at step 144 which identifies at least the dependent resources that are not cached locally at the mobile client 10 and the dependent resources that are cached locally at the mobile client 10 but have expired. The mobile proxy software 14 sends the HTTP response having the primary resource to the browser 12 at step 146 and transmits the acknowledgment list over the narrow-band connection to the fixed server 30 at step 148. Following step 148, the operation of the mobile client 10 again pauses while further steps are performed at the fixed server 30 as shown in FIG. 6.

Returning to FIG. 6, after the fixed server 30 receives the acknowledgment list from the mobile proxy software 14 (at step 150), the fixed proxy software 32 determines whether there are any dependent resources to be sent based upon the acknowledgment list and builds a summary response at step 152. The summary response comprises one or more status codes (51, 52 and 53 of FIG. 2), there being one status code for each dependent resource (54, 55 and 56 of FIG. 2) requested from the mobile client 10 in the acknowledgment list. If all of the dependent resources are retrieved by the fixed server 30, there will be one dependent resource for each status code in the summary response. If there are dependent resources that are not retrieved successfully by the fixed server 30, the status code corresponds to an error condition (discussed in more detail below) which informs the mobile client 10 not to expect those particular resources. The summary response may be a single transmission (as shown at step 154) or the summary response may be fragmented into several transmissions with the status codes for all dependent resources included in the first fragment. If the summary response is fragmented, the mobile client does not have to transmit a reverse channel acknowledgment across the narrow-band communications link for the individual fragments. Thus, the status codes located in the summary response corresponds to those resources, if any, that are to follow.

The mobile client 10 receives from the fixed server 30 the summary response and the dependent resources (if any) at step 156 of FIG. 4. Using the same information received from the fixed server 30 regarding the dependent resources in the summary response, the mobile proxy software 14 modifies its local cache 26 to prepare to respond to any forthcoming requests from the browser 12 at step 158. Once the cache 26 is prepared, the mobile proxy software 14 responds to the initial request from the browser 12, using the primary resource (either from its cache or from the fixed server's 30 compressed HTTP response) with an HTTP response containing the resource identified by the URL which the mobile user requested. When the browser 12 issues a request for any dependent resources, the mobile proxy software 14 is able to respond immediately to the request, or the mobile proxy software 14 is able to hold the request until the resource is received from the fixed server 30. As the dependent resources are received by the mobile proxy software 14, the mobile proxy software 14 updates its local cache and fulfills any requests from the browser 12 which have been held.

FIG. 7 is a flow chart of a method of operation for a cache check according to the invention. After a proxy examines its cache for a resource at step 180, the proxy must determine whether the requested resource is present in the cache at step 182 (i.e., whether the proxy has previously received the requested resource). If the resource is not present, the resource is not cached locally. If the resource is present, the proxy must check the expiration of the resource at step 184. Checking the expiration of the resource assures the proxy that the resource that is cached is up-to-date within a certain time frame (e.g., within 24 hours, etc. (depending on the nature of the resource)). If the resource has expired, the proxy must seek the resource from another source. If the resource has not expired, the proxy retrieves a dependency list at step 188. The proxy determines whether there are any dependent resources at step 190. If there are dependent resources, the proxy checks the cache further for the dependent resources at step 192. If the dependent resources are found in the cache, the proxy checks the expiration of the dependent resources at step 194. If the dependent resources have not expired, then they are valid. If the dependent resources have expired, then the proxy must seek the dependent resources from another source.

Standard HTTP requests and responses consists of a request or status line, zero or more headers consisting of a "field-name", a value and (optionally) an entity body. The request or status line, and the headers are ASCII text, separated by carriage-returns and line-feed control characters. The header (the request/ status line and the headers, collectively) is always transmitted ad the entity body, if present, is transmitted across the narrow-band communications links uncompressed. The wireless Web proxy protocol replaces the standard HTTP requests and responses with a binary format consisting of tokens for the standard parts of the request/ status lines and for the standard header "field-names" and common values. Non-standard field-names (e.g., "X-" headers) or values are left untouched.

Tokens are fixed predetermined elements of the wireless Web proxy protocol. Each proxy is knowledgeable of the information that is tokenized and its corresponding token. The use of tokens allows the "sender" proxy to transmit less data across the narrow-band communications link.

Further, request headers and response headers are cached at the proxies. Examples of header fields are content type, content length, content coding, character sets, etc. Having the request and response headers cached at the proxies allows the "sender" proxy to only send new or changed fields in the header across the narrow-band connection link to the "recipient" proxy.

Moreover, certain responses containing "dynamic" HTML form consist of a large amount of boiler-plate language and only a small amount of resource-specific information (e.g., the result of a search against a search engine or database such as an on-line phone directory or a stock quote service). Boiler-plate language is specific to a HTML page (e.g., the body of a response). Thus, with a large boiler-plate language, the possibility exists that the user will experience long delays to receive only a small portion of resource-specific information. In order to prevent the user from experiencing such long delays, the wireless Web proxy system caches the response at the respective proxies. When the "recipient" proxy request the cached response again, the "sender" proxy compares the cached response with a current response (e.g., a response retrieved from the origin host). The "sender" proxy identifies the boiler-plate language between the cached response and the current response and only transmit the information that is not cached at the "recipient" proxy over the narrow-band connection. The "recipient" proxy combines the cached information with the information that is received over the narrow-band connection to re-construct the complete dynamic response.

This method is also very useful for responses corresponding to error conditions. Normally, a response corresponding to a error condition consists of a status-line which includes a status code, a reason phrase, a protocol version, zero or more headers and an entity body. These elements are essentially static and provide no information beyond the status code itself, although they usually total to several tens or hundred of bytes. As a result, the possibility exists that the user will experience long delays because all of the elements are transmitted across the narrow-band communications link. The wireless Web proxy system prevents the user from experiencing long delays by caching the above mentioned elements at the proxies and transmitting only the status code, corresponding to the error condition, across the narrow-band communications link. The complete HTTP response is constructed, based on the status code, at the mobile client and sent to the browser.

Moreover, as shown in FIG. 2, T1 is the time between the initial HTTP request and the time when the first response (e.g., the primary response) is transmitted from the fixed server 30 to the mobile client 10. Having a short duration of T1 allows the browser 12 to display the general information to the mobile user in a short period of time. Allowing the mobile user to quickly gain access of general information gives the mobile user a opportunity to cancel the request or submit a different request before completion of the page if the mobile user does not desire to see the entire page (e.g., if the mobile user does not like the type of information displayed, if the mobile user can obtain what he/she is looking for by looking at the general information, etc.).

FIG. 8 is a bounce diagram of a alternative embodiment of the invention. The alternative embodiment is similar to the preferred embodiment. However, the alternative embodiment does not include the mobile proxy software 14 generating and transmitting an acknowledgment list to the fixed server 30 as does the preferred embodiment. Instead, the alternative embodiment allows the fixed proxy server 22 to transmit to the mobile client 10 a HTTP response having a page and dependency (HTML) list after the fixed proxy software 32 retrieves the primary resource and any dependent resources. In response to the mobile proxy software 14 sending the HTTP response to the browser 12, the browser 12 submits a HTTP request for the dependent resources to the mobile proxy software 14. The fixed server 30 transmits the dependent resources to the mobile client 10. The mobile proxy software 14 caches the dependent resources in its local cache 26 and sends the dependent resources to the browser 12 upon request.

With a single request, a primary resource, an acknowledgment list and a stream of dependent resources in a single transmission, the information necessary to fully render a complete page of information, for example a Web page, can be delivered to the mobile user's Web browser 12. While the invention has been described in conjunction with a specific embodiment thereof, it is evident that many alterations, modifications and variations will be apparent to those skilled in the art in light of the foregoing description. Thus, it should be understood that the invention is not limited by the foregoing description, but embraces all such alterations, modifications and variations in accordance with the spirit and scope of the appended claims.

Claims

1. A method of efficient delivery of information, performed at a fixed server, comprising:

   receiving a HyperText Transport Protocol (HTTP) request for a primary resource as identified by a uniform resource locator (URL);
   connecting to a host identified in the URL;
   receiving the primary resource from the host;
   inspecting the primary resource to identity dependent resources;
   pre-fetching and assembling the dependent resources; and
   forwarding the primary resource to a requester.

2. The method of claim 1 further comprising:

   caching the dependent resources in a local cache;
   waiting for an acknowledgment list from a mobile client requesting the dependent resources; and
   upon receipt of the acknowledgment list, transmitting the dependent resources that were requested by the mobile client to the mobile client.

3. A wireless network proxy having a fixed part and a mobile part, the fixed part comprising:

   a first digital processor comprising a resource transceiver and a summary response builder coupled to the resource transceiver, the summary response builder having an output;
   a first radio transceiver coupled to the output of the summary response builder for sending at least a summary response over a communications link to a browser;

   the mobile part comprising:

   the browser;
   a second radio transceiver;
   a cache; and
   a second digital processor having a resource transceiver and an acknowledgment list generator coupled to the cache and the resource transceiver and coupled to the second radio transceiver, the acknowledgment list generator having a comparator comparing received resources with cached resources.

4. A method of providing resources to a browser comprising:

   sending a HyperText Transport Protocol (HTTP) request from the browser to a mobile client;
   sending the HTTP request from the mobile client to a fixed server;
   sending the HTTP request from the fixed server to an origin host;
   receiving a HTTP response at the fixed server from the origin host;
   sending from the fixed server to the mobile client a list of resources;
   at the mobile client, comparing the list of resources with resources stored in a cache;
   sending from the mobile client to the fixed server an acknowledgment list selectively indicating resources from the list of resources;
   assembling, at the fixed server, resources selectively indicated by the acknowledgment list; and
   sending the resources selectively indicated by

the acknowledgment list from the fixed server to the mobile client in a single transmission.

5. The method of claim 4 further comprising the step of sending each resource from the mobile client to the browser upon request.

6. A method comprising:
at a fixed server:

receiving a HyperText Transfer Protocol (HTTP) response from an origin host;
identifying that the HTTP response corresponds to an error condition;
sending a status code to a mobile client; and

at the mobile client

constructing a complete HTTP response, based on the status code, for delivery to a browser.

7. A method of efficient delivery of information, such method comprising:

transmitting from a mobile client a request header in its entirety one time to a fixed server;
caching the request header at a cache on the fixed server;
for subsequent transmissions of information having the request header, transmitting information, different than that contained in the request header that is cached at the fixed server, from the mobile client to the fixed server;
transmitting from the fixed server a response header in its entirety one time to the mobile client;
caching the response header at a cache on the mobile client, and
for subsequent transmissions of information having the response header, transmitting information, different than that contained in the response header that is cached at the mobile client, from the fixed server to the mobile client.

8. A method of reducing data transfer over a narrowband connection comprising:
at a mobile client:

receiving a HyperText Transport Protocol (HTTP) request for a resource from a browser;
examining a cache on the mobile client for the resource;
identifying that the mobile client has already received the resource and needs to request the resource again;
informing a fixed server that the mobile client has previously cached the resource;

at the fixed server:

examining the resource as currently cached on the fixed server;
requesting the resource from an origin host;
receiving a HTTP response having an updated resource from the origin host;
comparing the resource that is cached with the updated resource; and
sending information to the mobile client which when taken together with the resource that is currently cached at the mobile client constitutes the updated resource.

9. A method to initiate a transfer of data over a narrowband connection, such method comprising:
at a fixed server:

receiving a HyperText Transport Protocol (HTTP) request for a primary resource;
retrieving the primary resource from an origin host if the primary resource is not cached at the fixed server;
retrieving the primary resource from the origin host if the primary resource is cached at the fixed server but has expired;
caching the primary resource at a cache on the fixed server;
identifying dependent resources from the primary resource;
requesting from the origin host the dependent resources that are not cached locally at the fixed server;
requesting from the origin host the dependent resources that are cached at the fixed server if the dependent resources cached at the fixed server have expired;
caching the dependent resources; and
sending the primary resource to a mobile client.

10. The method of claim 9 further comprising:
at the mobile client:

caching the primary resource at a cached on the mobile client;
identifying dependent resources from the primary resource;
generating an acknowledgment list identifying at least the dependent resources that are not cached locally at the mobile client and the dependent resources that are cached at the mobile client but have expired;
sending the acknowledgment list to the fixed server;
receiving the dependent resources that are not

cached locally at the mobile client and the dependent resource that are cached at the mobile client but have expired;
caching the dependent resources at the cache on the mobile client;
sending the primary resource to a browser; and sending the dependent resources to the browser upon request.

11. A method at a mobile client comprising:

receiving a request for a primary resource from a browser;
checking a cache on the mobile client for the primary resource;

if the primary resource is present at the mobile client:

(i) sending the primary resource to the browser;

if the primary resource is present ad has expired at the mobile client:

(i) tokenizing and compressing the request for the primary resource;
(ii) transmitting the request for the primary resource to a fixed server;
(iii) receiving different, updated information for the primary resource from the fixed server;
(iv) decompressing the different, updated information for the primary resource;
(v) assembling the different, updated information received from the fixed server with information previously cached at the mobile client for the primary resource;
(vi) updating the cache at the mobile client with a complete updated primary resource; and
(vii) sending the complete updated primary resource to the browser;

if the primary resource is not present at the mobile client:

(i) tokenizing and compressing the request for the primary resource;
(ii) transmitting the request for the primary resource to a fixed server;
(iii) receiving the complete updated primary resource from the fixed server;
(iv) updating the cache at the mobile client with the complete updated primary resource; and
(v) sending the complete updated primary resource to the browser.

12. The method of claim 11 further comprising, following a step of sending the complete updated primary resource to the browser, the steps of:
at the mobile client:

identifying dependent resources;
generating a acknowledgment list;
transmitting the acknowledgment list to the fixed server;
receiving a summary response followed by dependent resources from the fixed server;
updating the cache at the mobile client with the dependent resources; and
calculating a plurality of dependent resources.

13. A method at a fixed server comprising the steps of:

receiving from a mobile client a request for a primary resource;
checking a cache on the fixed server for the primary resource;

if the primary resource is present and valid at the fixed server:

(i) tokenizing and compressing the primary resource; and
(ii) transmitting the primary resource to the mobile client;

if the primary resource is present and stale at the fixed server:

(i) retrieving the primary resource from an origin host;
(ii) caching the primary resource at the fixed server;
(iii) tokenizing and compressing the primary resource; and
(iv) transmitting the primary resource to the mobile client,

if the primary resource is not present at the fixed server:

(i) retrieving the primary resource from a origin host;
(ii) caching the primary resource at the fixed server;
(iii) tokenizing and compressing the primary resource; and
(iv) transmitting the primary resource to the mobile client.

14. The method of claim 13 further comprising:

pre-fetching dependent resources for the primary resource; and
caching the dependent resources for the primary resource at the fixed server.

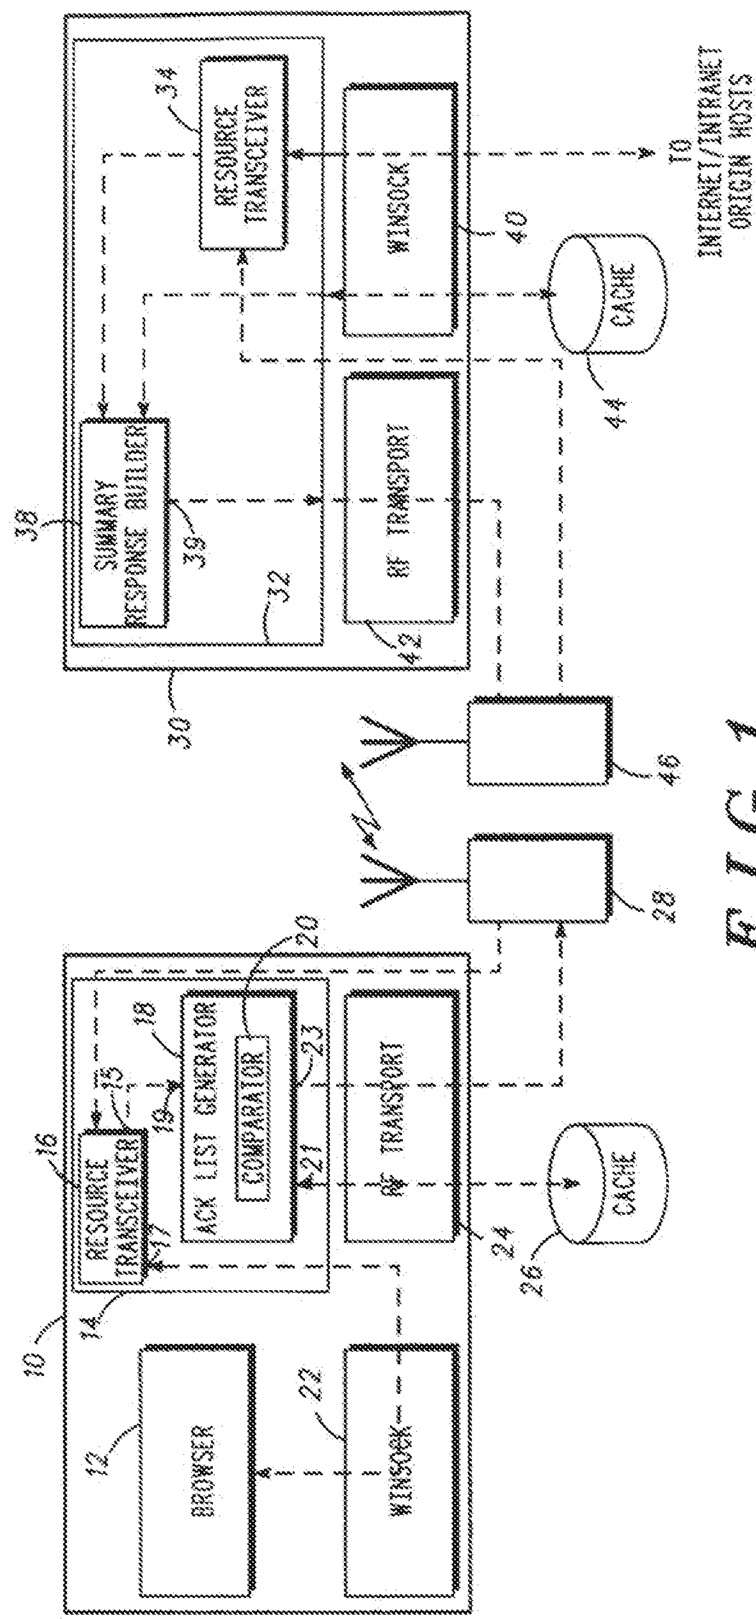
BROWSER
WINSOCK
RESOURCE TRANSCEIVER
ACK LIST GENERATOR
COMPARATOR
RF TRANSPORT
CACHE
SUMMARY RESPONSE BUILDER
RF TRANSPORT
WINSOCK
RESOURCE TRANSCEIVER
CACHE
TO INTERNET/INTRANET ORIGIN HOSTS

FIG.1

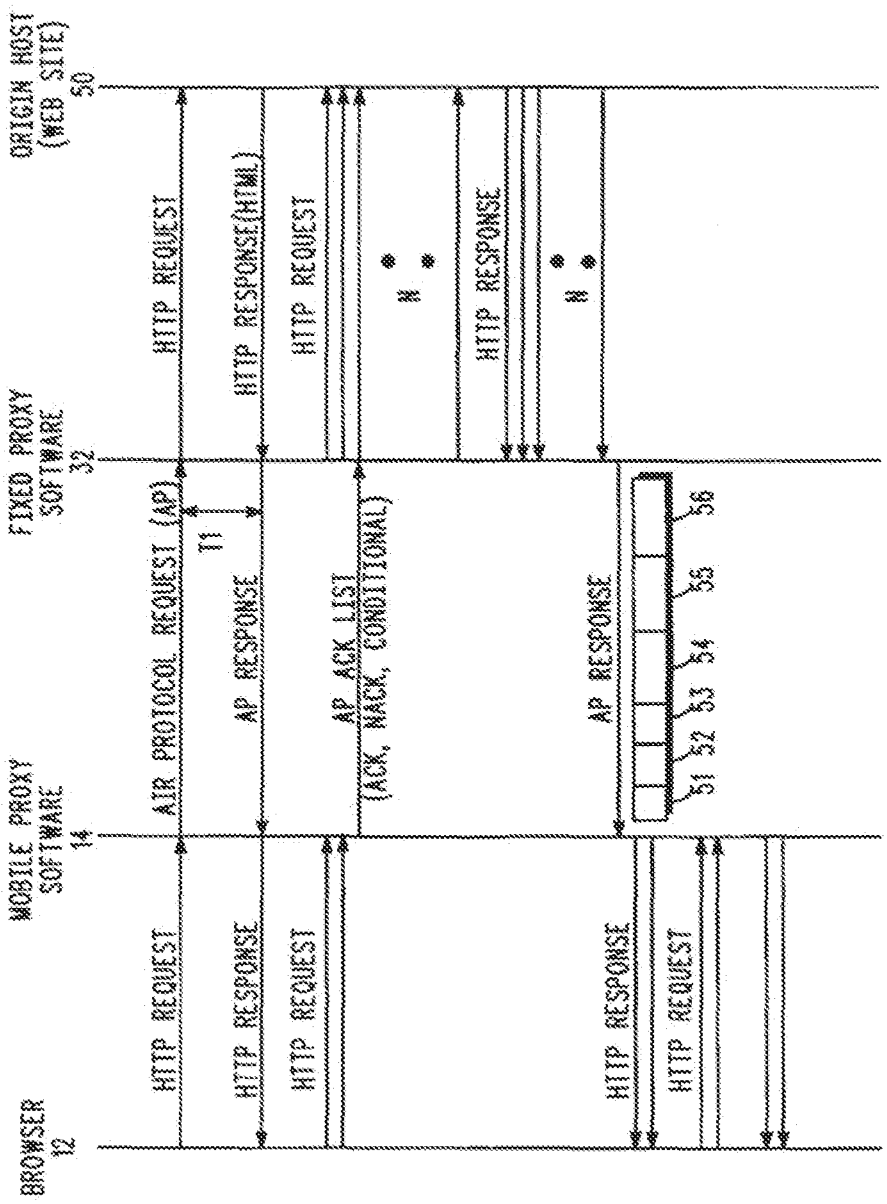
BROWSER
12
MOBILE PROXY
SOFTWARE
14
FIXED PROXY
SOFTWARE
32
ORIGIN HOST
(WEB SITE)
50
HTTP REQUEST
HTTP RESPONSE
HTTP REQUEST
HTTP RESPONSE
HTTP REQUEST
AIR PROTOCOL REQUEST (AP)
T1
AP RESPONSE
AP ACK LIST
(ACK, NACK, CONDITIONAL)
AP RESPONSE
51
52
53
54
55
56
HTTP REQUEST
HTTP RESPONSE(HTML)
HTTP REQUEST
N
HTTP RESPONSE
N

FIG. 2

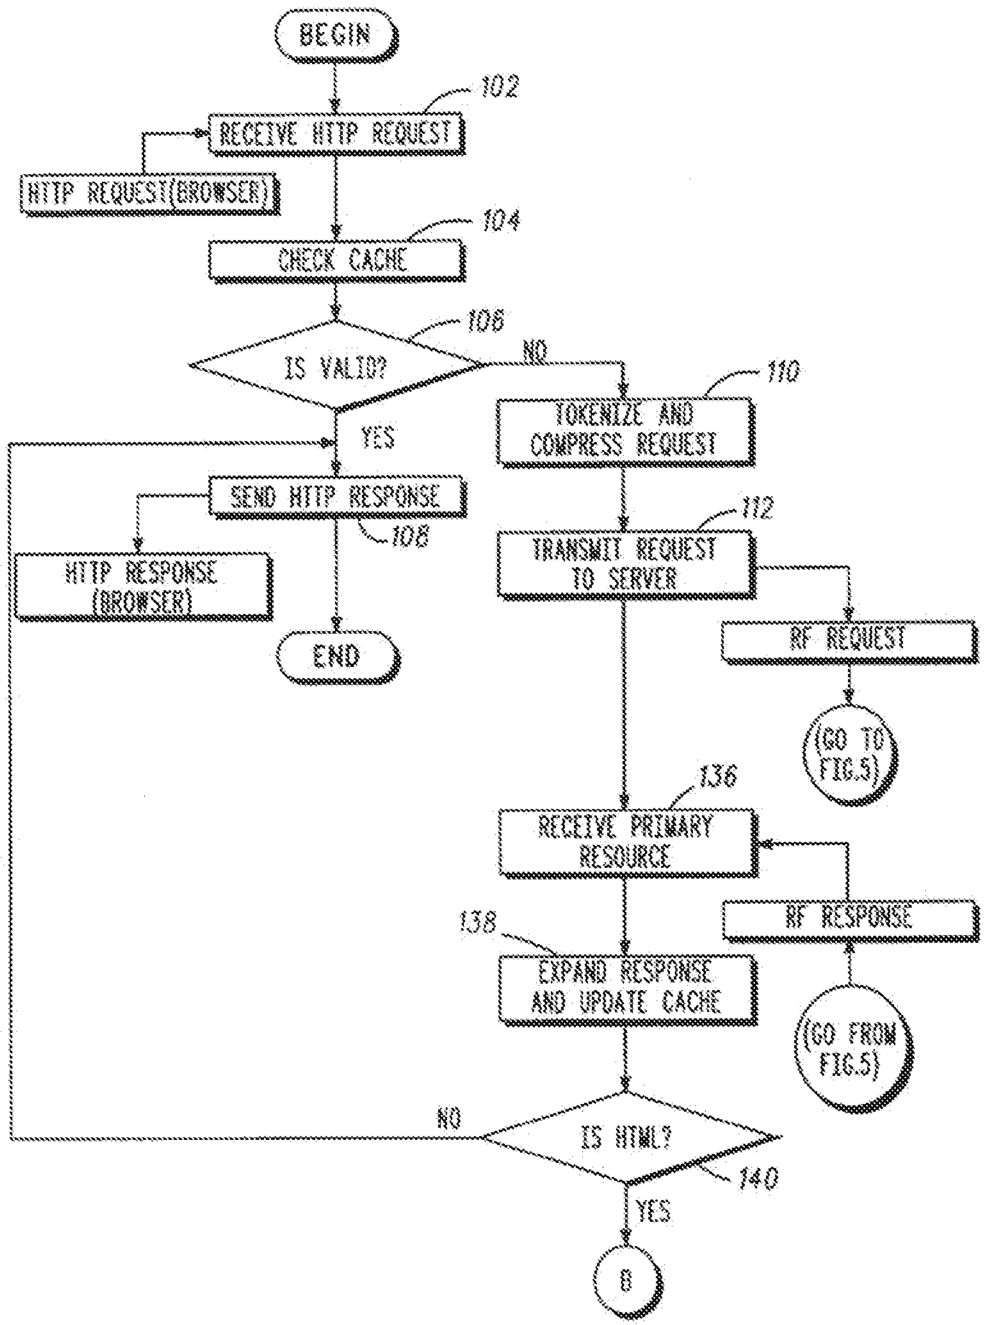
BEGIN
102
RECEIVE HTTP REQUEST
HTTP REQUEST(BROWSER)
104
CHECK CACHE
106
IS VALID?
NO
YES
110
TOKENIZE AND COMPRESS REQUEST
SEND HTTP RESPONSE
108
HTTP RESPONSE (BROWSER)
112
TRANSMIT REQUEST TO SERVER
END
RF REQUEST
(GO TO FIG.5)
136
RECEIVE PRIMARY RESOURCE
RF RESPONSE
138
EXPAND RESPONSE AND UPDATE CACHE
(GO FROM FIG.5)
NO
IS HTML?
140
YES
B

FIG. 3

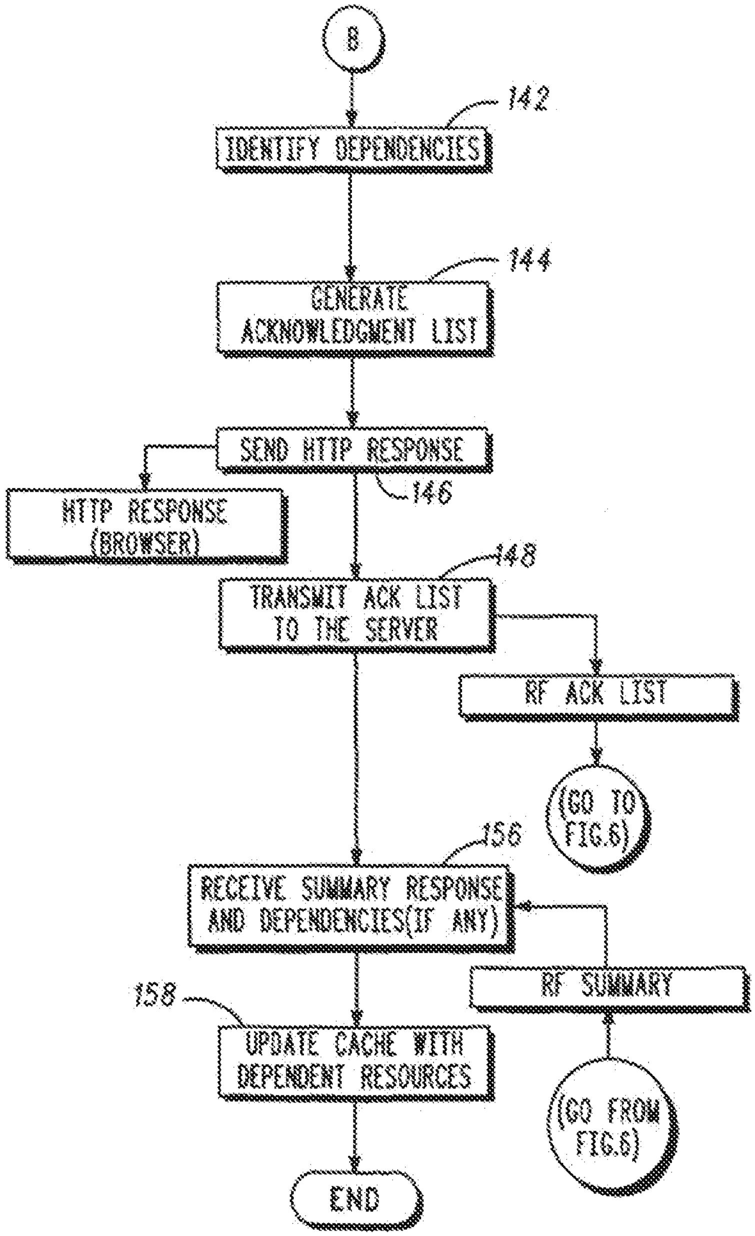
B
142
IDENTIFY DEPENDENCIES
144
GENERATE ACKNOWLEDGMENT LIST
SEND HTTP RESPONSE
HTTP RESPONSE (BROWSER)
146
148
TRANSMIT ACK LIST TO THE SERVER
RF ACK LIST
(GO TO FIG.6)
156
RECEIVE SUMMARY RESPONSE AND DEPENDENCIES(IF ANY)
RF SUMMARY
158
UPDATE CACHE WITH DEPENDENT RESOURCES
(GO FROM FIG.6)
END

FIG. 4

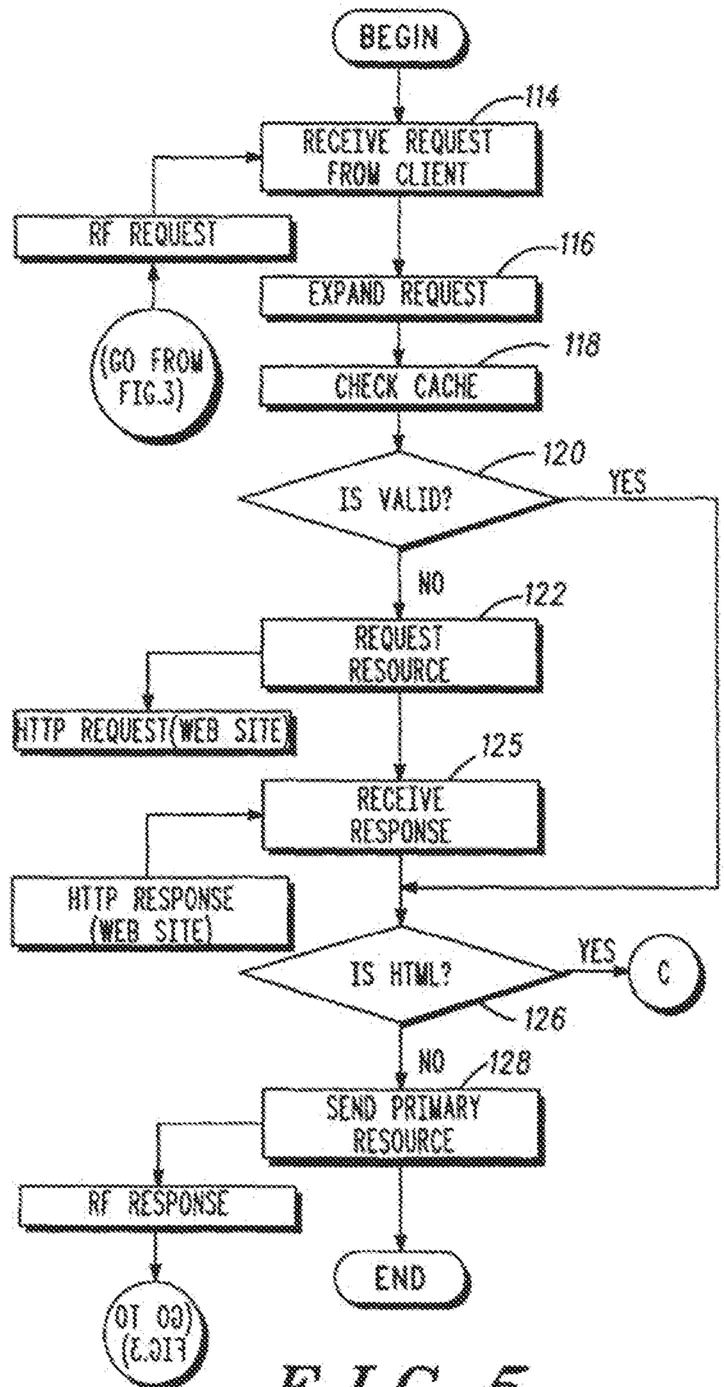
BEGIN
114
RECEIVE REQUEST FROM CLIENT
RF REQUEST
(GO FROM FIG.3)
116
EXPAND REQUEST
118
CHECK CACHE
120
IS VALID?
YES
NO
122
REQUEST RESOURCE
HTTP REQUEST(WEB SITE)
125
RECEIVE RESPONSE
HTTP RESPONSE (WEB SITE)
IS HTML?
YES
C
126
NO
128
SEND PRIMARY RESOURCE
RF RESPONSE
END
(GO TO FIG.3)

FIG. 5

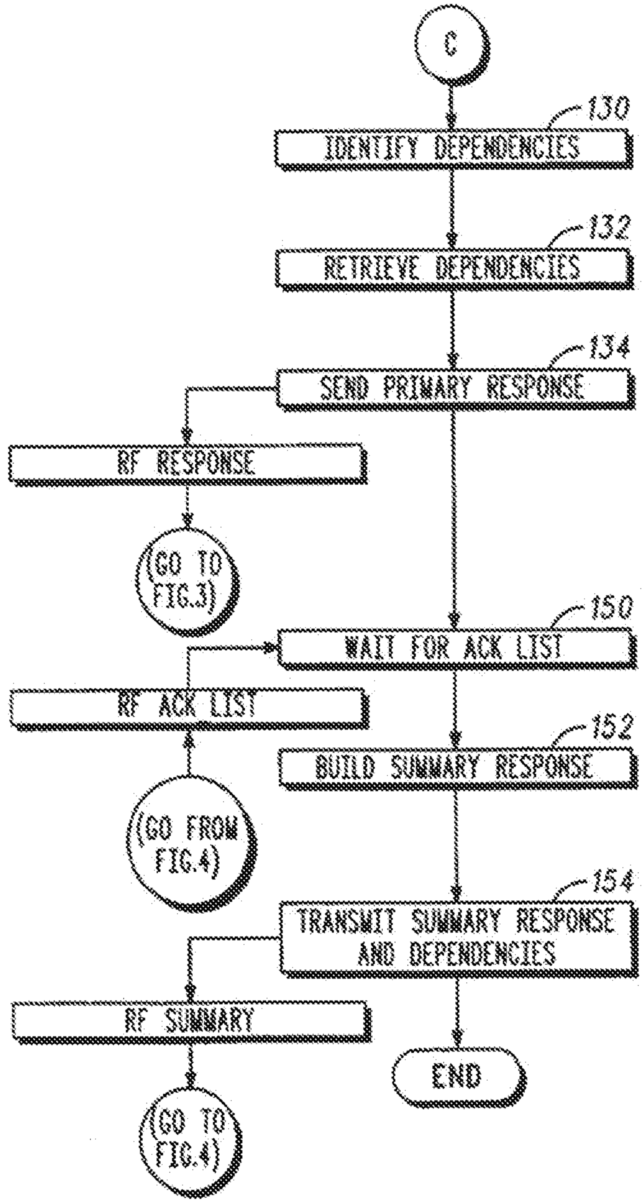
C
130
IDENTIFY DEPENDENCIES
132
RETRIEVE DEPENDENCIES
134
SEND PRIMARY RESPONSE
RF RESPONSE
(GO TO FIG.3)
150
WAIT FOR ACK LIST
RF ACK LIST
152
BUILD SUMMARY RESPONSE
(GO FROM FIG.4)
154
TRANSMIT SUMMARY RESPONSE AND DEPENDENCIES
RF SUMMARY
END
(GO TO FIG.4)

FIG.6

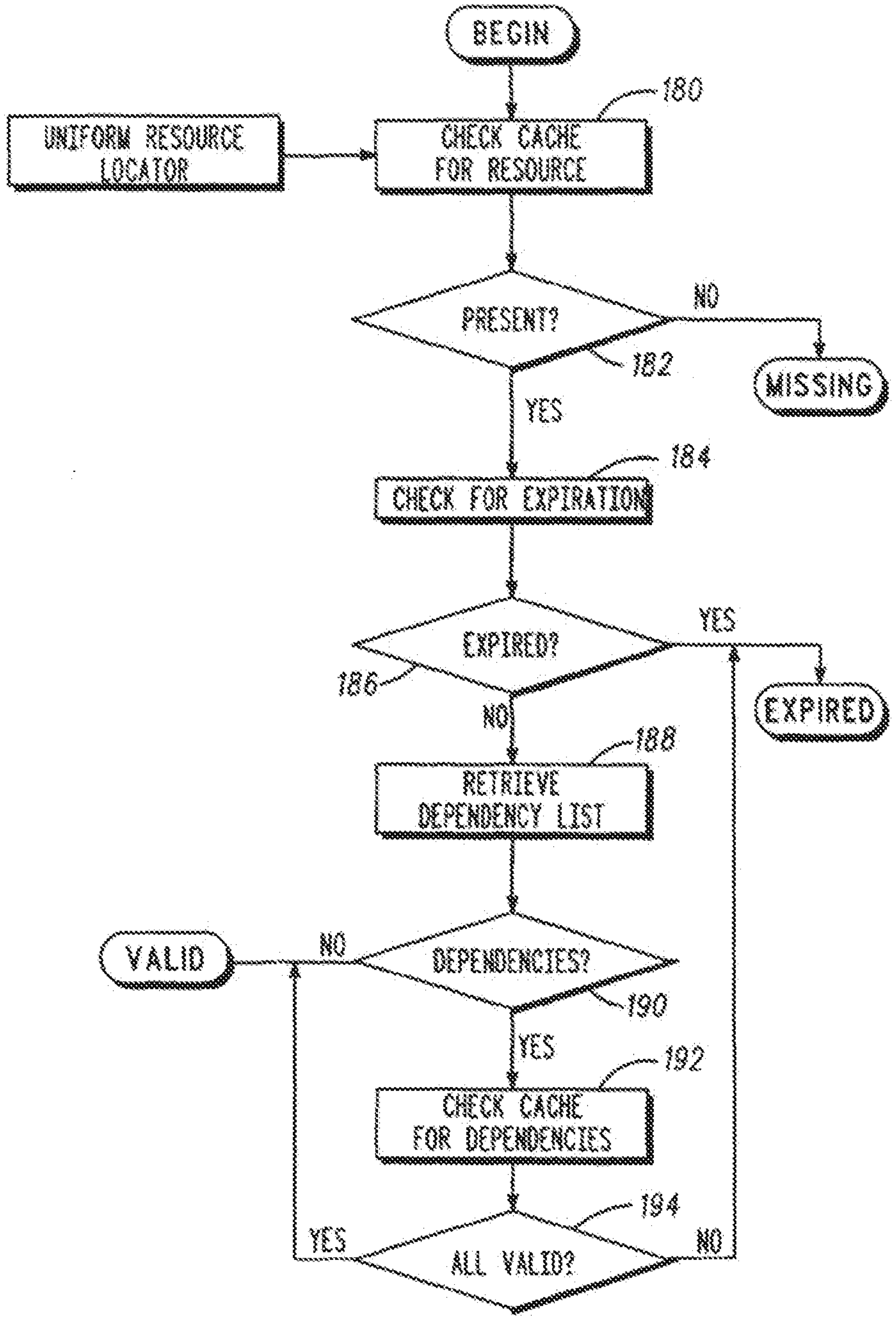
BEGIN
180
UNIFORM RESOURCE LOCATOR
CHECK CACHE FOR RESOURCE
PRESENT?
NO
182
MISSING
YES
184
CHECK FOR EXPIRATION
EXPIRED?
YES
186
EXPIRED
NO
188
RETRIEVE DEPENDENCY LIST
VALID
NO
DEPENDENCIES?
190
YES
192
CHECK CACHE FOR DEPENDENCIES
194
YES
ALL VALID?
NO

FIG.7

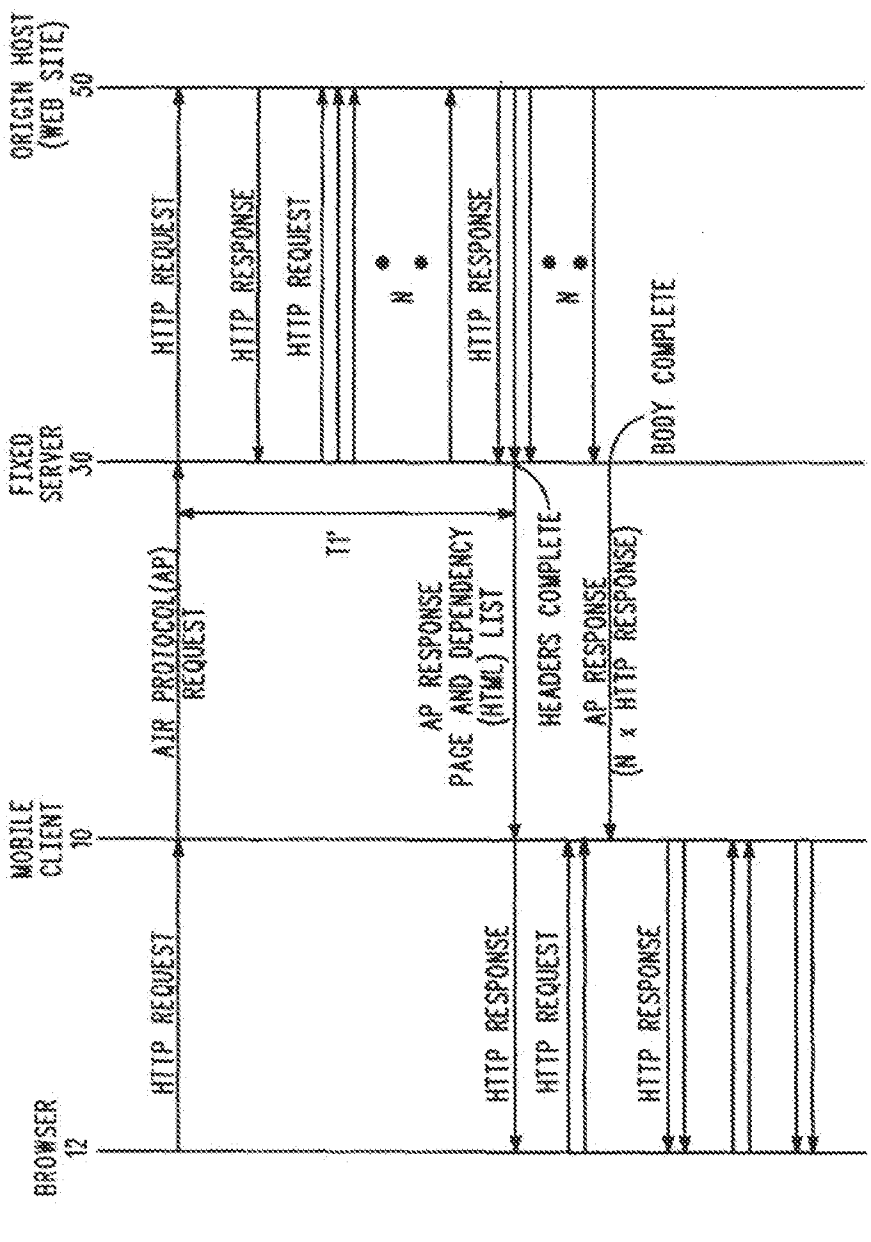
BROWSER
12
MOBILE
CLIENT
10
FIXED
SERVER
30
ORIGIN HOST
(WEB SITE)
50
HTTP REQUEST
AIR PROTOCOL(AP)
REQUEST
HTTP REQUEST
HTTP RESPONSE
HTTP REQUEST
TT
N
AP RESPONSE
PAGE AND DEPENDENCY
(HTML) LIST
HTTP RESPONSE
HEADERS COMPLETE
HTTP RESPONSE
HTTP REQUEST
N
AP RESPONSE
(N x HTTP RESPONSE)
BODY COMPLETE
HTTP RESPONSE

FIG. 8

(19) Europäisches Patentamt
European Patent Office
Office européen des brevets

(11) EP 0 875 823 A2

(12) **EUROPEAN PATENT APPLICATION**

(43) Date of publication:
04.11.1998 Bulletin 1998/45

(51) Int. Cl.[6]: **G06F 9/44**

(21) Application number: **97121695.7**

(22) Date of filing: **09.12.1997**

(84) Designated Contracting States:
AT BE CH DE DK ES FI FR GB GR IE IT LI LU MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: **24.04.1997 US 842460**

(71) Applicant:
Hewlett-Packard Company
Palo Alto, California 94304 (US)

(72) Inventors:
- **Wang, Lanzhong**
  Fort Collins, CO 80525 (US)
- **Cairns, Charles W. Jr.**
  Fort Collins, CO 80526 (US)

(74) Representative:
Schoppe, Fritz, Dipl.-Ing.
Schoppe & Zimmermann
Patentanwälte
Postfach 71 08 67
81458 München (DE)

(54) **Method and structure for tokenized message logging system**

(57) Methods and structure for storing log messages (400-410, 900-914) in a tokenized, international format and for presenting (viewing, printing, etc) (500-508, 600-608) the tokenized log message in a locally preferred native language. Log messages from a computing system (204) or application (202) are stored in an international tokenized format (206-210) which remains constant regardless of the particular nation in which the system or application is operated. The tokenized format includes a message ID field which identifies a unique message and includes parameters values which are to replace variable portions, if any, of the identified message. A plurality of localized message catalog files (302, 304) are available to retrieve a localized native language string which corresponds to each tokenized message. The message ID field serves as an index to the localized catalog files. A viewer program (300) then retrieves a localized text string message from a selected message catalog file (302, 304) replacing any variable portions identified therein with parameter values supplied in the tokenized message. The formatted, localized native language message is then presented (e.g., displayed, printed, etc. 308) to the requesting user in their preferred local native language. The present invention thereby enables users having a first preferred language (e.g., 100-104) to review or otherwise process logged messages generated by a system or application operable in an environment having a second preferred language (e.g., 110-114). The present invention also simplifies parsing problems relating to automated analysis and post-processing of logged messages. The tokenized international format of the present invention remains constant regardless of the preferred local language of the system or application which generated the message.

FIG. 2

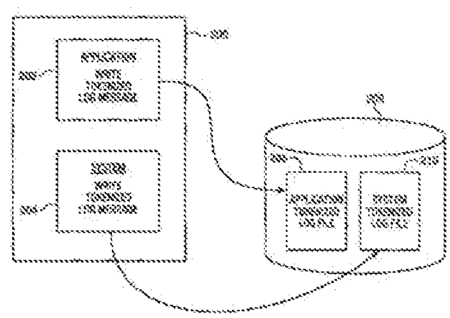

EP 0 875 823 A2

Description

Background of the Invention

1. Field of the Invention

The present invention relates to computing systems and application which generate log files and in particular to methods and structures for generating a parametarized, tokenized log file to ease multi-national language translation problems in multi-national computing environments.

2. Discussion of Related Art

Application of computing systems is growing rapidly and internationally. A single computing system or application, for example, may now need to concern itself with communicating to users in a plurality of languages. Or, for example, large networks of computing systems and associated distributed applications are commonly deployed internationally. Multi-national corporate or other entities may therefore require information generated in one locale of the distributed computing environment, using a first native language (e.g., English), to be understood by another user in a second locale using a second native language (e.g., French).

One common computing application which exemplifies such a multi-lingual requirement is the generation and viewing of log information within such computing systems and applications. It is common for example that an operating system maintains a variety of log files to allow a system manager to track certain activities, configurations, and performance. In like manner, application programs operable atop such operating systems may maintain their own log files to enable monitoring of that application's operations. Users of these log files may be geographically separated and linguistically diverse. For example, support engineers and technicians may review such log files to determine potential causes of computing problems related to the system or application. However, a specialist support technician (for example a regional or international expert in support of a particular system or application) may speak a different native language than the native language in which the log files are generated.

An undesirable solution to this multi-lingual computing problem is to simply adopt one language as a standard (e.g., English) and force users to manually (mentally) translate from the standard language to their preferred native language. For obvious reasons, this option is impractical.

Therefore, in a multi-lingual computing environment, designers and vendors of such operating systems and application programs need concern themselves with features which automate the translation of messages (including log messages) in to a local native language. This process is also referred to herein as "localization" of the system or program.

One approach to localization is to select a preferred local (native) language for the system or program at the time the system is installed and configured. The program designers therefore design their systems and applications such that messages generated thereby are adapted to the local native language at the time of installation. For example, a program or system installed and configured in Japan will be configured to generate messages (e.g., messages in log files) in Japanese while the same program installed in China will be configured to log messages in Chinese. In such systems, there may exist a plurality of message catalogs, one for each language. However, one such language is selected for configuration (localization) of each system. Log files are therefore written and viewed (or otherwise utilized) in the locally preferred native language of the system which generated the log messages. Specifically, a system or application generates a log message by looking up a desired message in the selected local catalog using a message index value. The system or application then writes the log message found in the catalog into the log file. Messages in the log file are therefore written in the locally preferred native language and are viewed or otherwise utilized in that local language.

However, if systems in a distributed computing environment select different localized configurations of the system or program, other nodes in the distributed environment will be unable to view the remote systems' log files. For example, an attempt by a user to view a log file from a system localized for Japanese on a system localized for Chinese will likely result in garbage information being displayed. The Japanese encoding of log messages does not correspond to the Chinese viewing capability of the local system.

A different aspect to program and system localization has been to utilize a large character code set (e.g., ISO 10646 Unicode) which includes characters for several local languages. Each local language is encoded as a subset of characters in such an extended character set. A local viewer of a remotely generated log file therefore has knowledge of which local language the log message was written in. This solution does not help the user perform the needed translation but rather serves merely to help the local user identify the translation required for the message. For example, a user of a Japanese localized system viewing a log message from a Chinese localized system will recognize that the log message is encoded from the Chinese subset of the internationally extended character set. With an appropriate viewer program , the Japanese user may view the remotely generated log message in its native language (e.g., Chinese) rather than viewing the remote message as garbage characters on the Japanese localized system. However, this approach does nothing to translate the remotely logged message into the preferred local native language (e.g., Japanese). The user is still required to manually translate the remotely logged message.

It is clear from the above discussion that it remains a problem to permit local review of remotely generated log messages in a locally preferred native language regardless of the remotely preferred native language.

Summary of the Invention

The present invention solves the above and other problems, thereby advancing the state of the useful arts, by providing methods and structures for generating and viewing log messages using an international tokenized log message format. This tokenized format is locally viewed by use of a viewer program in the preferred local native language. In particular, tokenized log messages generated in accordance with the present invention include a message ID value along with parameter values and control values used to add variable information to the message identified by the message ID. The tokenized message so generated is written to the system or application log file. The messages in the log file therefore are international by nature in that they contain little or no information which is unique to the specific local native language. Rather, the entire message is encoded by the message ID in combination with any specified parameter values.

The message ID value serves as an index into a local message catalog file. Each message in the local message catalog file is encoded in the locally preferred native language of the local system. Variable portions of the message are encoded so as to identify (positionally) which of the parameter values of the logged message are to be used when viewing the logged message. A viewer program therefore operates on the local system at which a user wishes to view the logged message. The message viewer program retrieves a tokenized log message (e.g., from a locally or remotely generated log file) and presents it to the user in the users locally preferred native language. First the message ID of the tokenized log message is used to locate and retrieve a corresponding message (a message encoded in the locally preferred native language) in a selected message catalog. Next, any variable portions encoded in the retrieved local message are replaced by a corresponding parameter value supplied in the tokenized log message. The correspondence of variables and parameter values is preferably positional so that positional difference among local languages may be represented in the local message catalogs. Finally, the tokenized message translated to the selected native language is presented to the user (e.g., displayed, printed or filed as desired for the particular viewer program).

The present invention therefore solves problems of prior approaches by writing (generating) log messages and storing them in an international format. Each user wishing to view (or otherwise utilize) the international, tokenized log messages may do so with a viewer program using a localized native language message catalog for the translation process. For example, a tokenized log message file generated by a system or application localized for Chinese language operation may be viewed (or otherwise utilized) on another system localized for other than Chinese local language operation (e.g., English, Japanese, etc.). For example, a Japanese support engineer may thereby review, in Japanese, a log file generated by a customers system or application localized for French language operation.

Standard C programming language library functions such as **catopen()**, **catclose()**, **catgets()**, and **catgetmsg()** are exemplary of well known API functions which may be useful in implementing a viewer program in accordance with the methods of the present invention. A localized message catalog may be accessed by such API functions to retrieve the template used to display (or otherwise present to a user) a tokenized message in the preferred local language and style. The catalog template entries may, for example, comprise C language **sprintf()** format strings used by the viewer program to format any supplied parameters and to generate an output string from the format (catalog template entry).

An additional benefit realized by the present invention is that the tokenized log message format may be more easily parsed by computer programs wishing to analyze or otherwise process the logged information. Automated parsing of logged information in accordance with prior techniques required the parsing program to be aware of linguistic aspects of the log file (i.e., the language the log file was generated in). Locating the position of a particular parameter value in a message logged in accordance with prior techniques required the parser to be aware of the language in which the message was generated due to sentence structure, vocabulary, and character set differences among several languages. Tokenized log messages in accordance with the present invention may be parsed independent of the local native language of the computing system. The message itself and associated parameter values are easily located in the tokenized log message regardless of the local native language of the generating computer system. The tokenized message format of the present invention is constant regardless of the preferred local language of the computer on which the message was generated and regardless of the preferred local language of the system on which the parsing is performed.

Additional advantages of the invention will be set forth in part in the description which follows, and in part will be obvious from the description, or may be learned by practice of the invention. The advantages of the invention may be realized and attained by means of instrumentalities and combinations particularly pointed out in the appended claims.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a block diagram of a distributed computing environment in which the methods and structures of the present invention may be advantageously applied;

FIG. 2 is a block diagram of a computing process operable in accordance with the present invention to generate a tokenized message log file;

FIG. 3 is a block diagram of a computing process operable in accordance with the present invention to retrieve tokenized log messages and to present the message in localized format;

FIG. 4 is a flowchart describing the operation of a method of the present invention to generate and store a tokenized log message as shown in FIG. 2;

FIG. 5 is a flowchart describing the operation of a method of the present invention to retrieve a previously generated tokenized log message and present the message in localized form as shown in FIG. 3;

FIG. 6 is a flowchart describing the operation of the parameter substitution processing performed as a step in the method of FIG. 5;

FIG. 7 is a graphical presentation of Japanese text of exemplary entries in a Japanese message catalog;

FIG. 8 is a graphical presentation of Japanese text of an exemplary presentation of a tokenized log message file presented in accordance with the exemplary catalog entries of FIG. 7; and

FIG. 9 is a flowchart describing the operations of a method of the present invention to generate a tokenized message using templates in a canonical message catalog.

DESCRIPTION OF THE PREFERRED EMBODIMENT

While the invention is susceptible to various modifications and alternative forms, a specific embodiment thereof has been shown by way of example in the drawings and will herein be described in detail. It should be understood, however, that it is not intended to limit the invention to the particular form disclosed, but on the contrary, the invention is to cover all modifications, equivalents, and alternatives falling within the spirit and scope of the invention as defined by the appended claims.

TOKENIZED MESSAGE APPLICATION

FIG. 1 is a block diagram of an exemplary multi-lingual distributed computing environment in which the methods and structures of the present invention may be advantageously applied. Systems 100, 102, and 104 are interconnected via local area network (LAN) 106. These systems may represent, for example, a single workgroup or location of a multi-lingual computing enterprise. These systems 100-104 share a first local native language (e.g., English). A second workgroup of systems 110-114 are interconnected via LAN 116 and share a second local native language (e.g., Japanese). Still a third workgroup of systems, 120-124, are interconnected via LAN 126 and share a third native language (e.g., Chinese). All systems are interconnected via wide area network (WAN) 130.

Systems in each workgroup are typically configured to operate in their respective local native language. As noted above with respect to prior techniques, prior logging methods logged information in the locally configured native language. However, it may be desirable for a user (e.g., support engineer, analyst, or manager) to view on his/her local system a log file generated from another system. In accordance with prior techniques, so long as the two systems (the viewing system and the log generating system) are configured for the same local native language, there are few if any problems encountered. Standard file sharing or networking techniques are used to retrieve a log file from another system and the log file is simply viewed on the users system in the same language in which it was generated. However, as noted above, where in accordance with prior techniques the log file is generated in a different native language than that of the viewing system, unexpected results may be obtained.

In accordance with the methods and structures of the present invention, a tokenized log file generated by any system in the enterprise may be viewed in any selected language (including the local native language of any system in the enterprise). For example, a tokenized log file generated on an English configured system 100 may be viewed in English by a viewer program on any of systems 100-104, in Japanese by a viewer program on any of systems 110-114, or in Chinese by a viewer program on any of systems 120-124. The tokenized log message format of he present invention is therefore also referred to herein as language independent. The translation of the tokenized (language independent) format to a local native language by a viewer program generates a local native language or language dependent presentation of the tokenized log messages.

FIG. 2 is a block diagram of a system 200 in which an application process 202 writes a tokenized application log file 208 on a storage device 206. In addition, a system process 204 writes a system tokenized log file 210 on storage device 206. One skilled in the art will recognize that the techniques of generating log files may be performed by a plurality of processes within a system. Each process may generate one or more such tokenized log files either by direct operations to write tokenized log messages to a log file or via centralized services provided within a system environment. A plurality of log files may be used to distinguish certain classes of log messages such as system/application, errors/information/warning, etc. In the alternative, a single, centralized, log service may generate a single log file wherein the log entries include tokenized information representing the class or type of log message.

FIG. 3 is a block diagram of a system 300 used to view (also referred to herein as present) a previously generated tokenized log file. A viewer process in system 300 receives a previously generated log file from storage device 206 (e.g., application tokenized log file 208 or system tokenized log file 210). Alternatively, a tokenized log file may be retrieved from a remote system via LAN 306. The viewer program is configured to view tokenized log files in a selected local native language. A plurality of localized message catalogs 304 are available to the system 300 on storage medium 302 (e.g., one for each possible local native language). As above, the source of the message catalog files may be any local storage device or remotely attached storage via well known network connections and protocols. The viewer process in system 300 then translates and transforms the language independent format of the tokenized log messages into localized messages for presentation to (viewing by) the user. The presentation of such localized log messages may be by display on a computer screen, printing, generated speech, or other well known user interface techniques.

TOKENIZED MESSAGE GENERATION AND PRESENTATION METHODS

FIG. 4 is a flowchart describing the operation of methods of the present invention to generate tokenized log messages in a tokenized log file. The method of FIG. 4 is operable in a system such as shown in FIG. 2 to generate tokenized log files. One skilled in the art will recognize that the method of FIG. 4 is operable at many levels including application programs as well as operating system processes. Similarly, it will be recognized that the method of FIG. 4 is operable to store the generated messages in a log file as well as other storage devices and formats. Further, the generated messages may be transferred via networks to a remote process for concurrent viewing in the preferred local native language of the remote site. In like manner, the method of FIG. 4 may send the tokenized messages to communication media for transmission to other processes for further manipulation, storage, or viewing. Well known interprocess communication between processes operable on the same computer or on distinct computers may be applied for such message communications.

Element 400 of FIG. 4 is first operable to generate a header portion of the tokenized log message (also referred to herein as a token header). As discussed below, the header portion includes a message ID field which identifies the particular message being generated in a language independent manner. The message ID field is preferably a structured field including a number of subcomponents to identify a particular catalog, a particular subset within the catalog, and a particular message within the subset. Such hierarchical definition of a message ID allows improved flexibility in the maintenance of the message catalogs. Other fields may be included in the header portion as discussed below to permit other parsing and analysis of the messages in the message log.

Some messages are static in nature while others are dynamic in that the message needs to be customized to indicate a specific value or state. Such dynamic messages are encoded with variable portions to be replaced by particular values when viewed by a user. For example, a status message may indicate the current utilization level of a resource. The tokenized message therefore includes the present value of the resource utilization. Such values are referred to herein as parameter values or token values.

Element 402 is next operable to determine if any parameter values (token values are required for the particular message being generated. The parameter values (token values), if any, are appended to the tokenized message under construction. If the message has no variable portions, no dynamic features, no parameter values are required. If element 402 determines that parameter values are required for the message being generated, element 404 is operable to generate the requisite parameter values as a parameter portion to be appended to the tokenized message. Otherwise, element 406 is operable to generate a null parameter portion to be appended to the tokenized message under construction.

In both cases, element 408 is next operable to append the generated parameter portion to the tokenized message under construction. Preferably, the parameter portion is appended to the message after the token header with a delimiter character separating the two portions. In the case of a null parameter portion, the delimiter character will be followed by no parameter values.

Lastly, element 410 is operable to write the tokenized message comprising the token header and parameter portion to a log file for persistent storage. As noted above, the tokenized log message may be written to any of several well known storage media either locally or remotely accessed. In the alternative, the tokenized log message may be written directly to other processes (either local or remote) for further processing using well known interprocess communication techniques.

FIG. 5 is a flowchart describing the method of the present invention to present a tokenized log file entry (a tokenized message) in localized, native language format. The method of FIG. 5 is operable in a system such as shown in FIG. 3 to present tokenized messages retrieved from log files (either local or remote). It will be recognized by those skilled in the art that the method of FIG. 5 is operable to retrieve the generated messages from a log file (local or remote) as well as other storage devices. In like manner, the method of FIG. 5 may receive the tokenized messages from communication media on which other processes (local or remote) have transmitted the tokenized log messages. Well known inter-process communication between processes operable on the same computer or on distinct computers may be applied for such message communications.

Element 500 of FIG. 5 is therefore operable to receive a tokenized message for presentation (viewing). As noted herein, the message may be retrieved from storage devices (either local or remote) or may be received directly from a generating process via a communication medium and well known inter-process communication techniques. Element 502 is next operable to locate and retrieve an entry in the selected local message catalog having a message ID corresponding to the message ID in the header portion (token header) of the tokenized message received (or retrieved) by operation of element 500. As noted above, the local system on which the method of FIG. 5 may be configured to operate in any of several locally preferred native languages. This configuration operates to select among a plurality of possible message catalogs. Other portions of the header portion (token header) may be used to further refine the selection of a particular message catalog file. For example, there may be separate catalogs identified by fields in the header portion for each of several application processes or for system processes. Further, particular subsets of particular catalogs may be identified by fields in the header portion of the tokenized message. Well known techniques for file naming and file indexing may be applied by operation of element 502 to locate and retrieve the message catalog entry corresponding to the received tokenized message.

The entry retrieved from the catalog is also referred to herein as a template in that it provides a template for presentation of the tokenized message in the preferred local native language. The template is therefore preferably copied to a working buffer for purposes of generating the intended localized message from the retrieved template. If the content of the message as presented in the native language is static, then the template retrieved from the selected catalog and copied to a working buffer is the entire message to be presented. If the message has dynamic aspects, then as noted above, parameter values in the parameter portion (token values) of the tokenized message will be substituted for designated variable portions of the template to generate the intended localized message for presentation. Element 504 is therefore operable to determine whether the template contains any variable portions to be replaced by parameter values (token values) from the parameter portion of the tokenized message. If such variable portions are present in the template, then element 506 is operable to substitute corresponding parameter values from the parameter portion of the tokenized message into the buffered copy of the template in place of variable portions.

Finally, element 508 is operable to present the tokenized message as completely converted to localized native language format. Any variable portions in the message have been replaced by corresponding parameter values and the message is now presentable to a user in the local native language. Presentation of the message may include, for example, printing the message, displaying the message on a display screen, generating voice audio data corresponding to the message, as well as other well known user interface methods for the presentation of information.

The methods of FIGS. 4 and 5 are presented as though a single message is generated and presented. Those skilled in the art will recognize that repeated application of such methods will generate and present a plurality of such log messages.

FIG. 6 is a flowchart describing the operation of element 506 of FIG. 5 in additional detail. Element 506 is operable to substitute token values (parameter values) from the parameter portion of the tokenized message into corresponding variable portions of the template used to construct the localized native language message. As noted, log messages may include dynamic portions which are replaced at time of generation with actual values appropriate to the message and generating system. Templates in the local message catalog files therefore include an encoding of variable portions associated with a template.

In a simple embodiment, each variable portion in a template from start to end is replaced by a corresponding parameter value in sequential order as they are provided in the parameter portion. In other words, the first variable portion in a template is replaced by the first parameter value in the parameter portion, the second by the second, etc. However, different native languages often requires alteration of the order of words or phrases in a particular sentence structure. A strict replacement of each variable portion in a template, in sequence, by the corresponding parameter value, in sequence does not allow for such order differences among native languages when a tokenized message is presented. In the preferred embodiment of the present invention, the variable portions of the template therefore include positional indicia indicative of which of the parameter values is to be used in the substitution.

For example, in first native language, a message may be structured as (where variable portion substitutions are indicated in ***bold italics*** and wherein the text would be presented in appropriate text for that language):

There are presently ***10*** processes running using ***100 MB*** of virtual memory. Whereas in a second language, the same tokenized message may be phrased as:

There are ***100 MB*** of virtual memory in use by ***10*** running processes. Strict sequential substitution of parameter values for the variable portions cannot handle such linguistic differences. The preferred embodiment would therefore allow the first template to specify a different order of parameter substitution than the second template. Such template encoding may be implemented by well known features of the C programming language. So-called format specifications in the **sprintf()** function calls permits the specification to determine which parameter is to be used in the substitution. Specifically, a standard format specification is used as noted above by applying to the next argument (parameter value) in sequence as supplied to the **sprintf()** function call. However, The printf() family of functions supports random access to arguments. Conversions can be applied to the nth argument in the argument list, rather than to the next unused argument. In this case, the conversion character (%) is replaced by the sequence %digit$, where digit is a decimal integer n in the range [1,{NL_ARGMAX}], giving the position of the argument in the argument list. This feature provides for the definition of format strings that select arguments in an order other than standard sequential order as provided in the **sprintf()** function call.

Therefore, in the preferred embodiment, message catalog template entries include (among other fields) a C programming language format specification field which describes the local native language message corresponding to a particular message ID. The format specification includes conversion fields for any variable portions of the message to define the format of the parameter value (e.g., numeric, character string, floating point, data, time, etc.) to be substituted for the variable portion of the template. The conversion fields use the random access format specification features described above to select the order of the parameter value substitution as required by the local native language.

FIG. 6 therefore describes the operation of element 506 to substitute parameter values for the variable portions of the template corresponding to the tokenized message ID. Element 600 is first operable to locate a next variable portion in the template copy. Element 602 then locates and retrieves the parameter value from the parameter portion of the tokenized message in accordance with the random access sequence value of the **sprintf()** conversion field defining the variable portion of the template. Element 604 then transforms the retrieved parameter value in accordance with the conversion specification defining the variable portion of the template. Element 606 replaces the variable portion of the template (the conversion field) with the parameter value as transformed by operation of element 604. Lastly, element 608 determines whether additional variable portions remain in the template. If more variable portions remain in the template, processing loops back to element 600 until all variable portions of the template (conversion fields) have been processed. In essence, element 506 may be implemented as a standard **sprintf()** function invocation. However, those skilled in the art will recognize other equivalent implementations which may be better suited to particular environments. Element 506 may therefore be implemented by any of several well known formatting techniques including standard C programming language library function calls. In particular, the preferred embodiment permits conversion specification which encode special parameters as described below. Such special conversion specifications are outside the scope of the C programming language standard library function call. It is therefore preferred in the best presently known mode of practicing the invention to utilize a customized conversion library function modeled substantially in accordance with the C language **sprintf()** library function.

SPECIAL PARAMETERS

As noted above, the conversion specifications permitted by the present include format conversion specification which define special parameters useful in encoding complex messages in multiple native languages. For example, the date and/or time at which a message was generated as indicated by the message header portion time stamp value may be presented by to a user by a viewer program. Or for example, a process ID of the process that generated the messages (as distinct from the viewer process which is presenting the tokenized message to the user) may be useful for presentation in a message. Special parameters features of the present invention provide a means to enter data into the parameter list of a tokenized message without requiring the programmer of the application generating the tokenized message to obtain and explicitly enter the value of that parameter. In addition, such special parameters may be automatically formatted in the locally preferred native style (e.g., date and time formats specific to a local language).

Standard C language style format specifications do not include a format conversion specification which would have special knowledge of such encoded parameters. It is therefore useful to have special format specifications available in encoding the template entry of the catalog which specifies that such special parameters be presented in the format preferred by the local native language. In general, format specifications customized to such special presentation needs of local native languages are encoded as special format specifications in the template string.

Special parameter values precede all other parameter values in the parameter portion of the tokenized message. As noted above, a value in the header portion of the tokenized message identifies the number of such special parameter values which are stored at the start of the parameter portion of the tokenized message.

A format specification of "%D" is used, for example, to indicate that the date (e.g., the time stamp date in the header portion) be displayed or printed in the format preferred by the local native language. Likewise, a "%T" format specification could be used to present a time value (e.g., the time stamp value in the header portion of the message) in a locally preferred format. Other special parameter types might include: a format specification of a process ID value (e.g., "%P"), a format specification for the last system error number (e.g., "%m"), the name of the process which generated the log file entry, the total memory used by a particular process, etc. Those skilled in the art will recognize a variety of such format conversion specification types for formatted presentation of equivalent special parameters. Such special format conversions are extensions and enhancements to the standard conversions commonly available in the C programming language.

MESSAGE NESTING SPECIAL PARAMETERS

Another type of special parameter encoding enables the nesting of messages within a message template. As noted above, each message is associated with a message ID value used as an index to locate the corresponding template in the selected catalog. A nested message is one which is used to complete the presentation of another message. In other words, a first template in a message catalog may include a special parameter the value of which is a message ID value which identifies another (a second) template to be inserted (nested) at that point in presentation of the first message template.

As a simple nested message example, consider the two sentences:

"The man opened the window."
"The man opened the door."

Using the format conversion specification "%Z" to indicate a nested message special parameter, a message template catalog might include the following templates:

Message ID	Message template string
12340	"the door"
12341	"the window"
12342	"The man opened %Z."

A tokenized log message might therefore include the following to represent the sentences:

"The man opened the door."
"The man opened the window."

Message ID	Date and other header information	Parameter values
12342	...	12340
12342	...	12341

In this simple example, the messages 12340 and 12341 are nested within message 12342. Multiple such embedded or nested messages may occur within a message and the nested message may itself include parameters (including special parameter values).

TOKENIZED MESSAGE DATA STRUCTURES, FORMATS, AND EXAMPLES

As noted above, a tokenized message comprises a header portion and a possible null parameter portion. The header portion identifies the message and the appropriate message catalog to be used in presenting it in localized form. One exemplary data structure which may be used to describe the header structure is as follows (in C-like pseudocode):

```
struct tok_msg_hdr {
   boolean flags[NFLAGS],          /* flags used to classify and analyze messages */
   struct msg_ID {
      id_type msg_identifier,      /* identifier value for message */
      set_type set_identifier,     /* subset of catalog in which message is found */
      cat_type catalog_identifier  /* catalog in which message is found */
   } ID,
   time_t time_satmp,              /* date/time of message generation */
   int s_parm_present              /* number of special parameters present */
}
```

Following the header portion of the tokenized message is the parameter portion where the parameter values (token values) are appended to the message. For example, character string values may be appended encoded between matched quotation marks, numeric values may be appended encoded in decimal digits, booleans may be encoded as decimal zeros or ones, etc. A variety of encoding techniques may be applied to the construction of the header portion and the parameter portion. The above pseudo-code data structure is therefore intended only to suggest the type of information which may be useful in conjunction with the methods of the present invention. Character encoding of the header portion may be preferred for ease of porting and human readability while binary encoding of various values may be preferred to reduce the size of tokenized log messages and log files.

As noted above, the message catalog entries (templates) are constructed to correspond to the message ID information encoded in tokenized message headers and map that ID to a sprintf() style format specification used to convert the tokenized message for localized presentation. As noted the format specification provides the text in local native language along with the conversion fields for any variable portions of the message. The conversion fields preferably utilize the random access feature to allow re-ordering of parameter values as required for a particular local native language sentence structure. As a pseudo-code data structure, a message catalog template entry may be encoded as follows:

```
struct tok_msg_hdr {
   struct msg_ID {
      id_type msg_identifier,      /* identifier value for message */
      set_type set_identifier,     /* subset of catalog in which message is found */
   } ID,
   char template[ ] =              /* exemplary template with variable portions */
        "There are %2$d processes running using %1$d of memory\n",
}
```

The cat_identifier portion of the ID is subsumed in the identification of the catalog which contains the template and is therefore removed from the ID field within the catalog entry. Otherwise, the ID field of the template is used to match the ID field of the tokenized message to be presented. Any of several well known search techniques and associated structures may be employed to speed the search for the proper catalog to localize a tokenized message as well as to speed the search for a particular template within the catalog.

The localization of language independent tokenized messages may be seen in the following examples. A sample tokenized log file (set of log messages) may include the following where "!" is a delimiter character to delimit the header portion from the parameter portion and wherein the parameter portion (if any) appears indented on the line below the header portion for typographic clarity herein:

0	0	140267	14	267	858123468	1	!
	P1237 "install" "AGENT SESSION" "swposix-0312"						
0	7	140068	14	68	858123468	1	!
	P1237 "Agent Session" "root@swposix.fc.hp.com"						
0	7	140024	14	24	858123526	0	!
	"swi18n.fc.hp.com" "/var/spool/sw"						
0	7	140025	14	25	858123526	0	!
	"swposix" "/"						
0	7	140033	14	33	858123567	0	!
0	11	140061	14	61	858123567	0	!
	" " " " "test_token.ok_filest,r="						
0	11	140065	14	65	858123567	0	!
	" " " " "test_token.skip_filest,r="						
0	7	140347	14	347	858123567	0	!
	1 2						
0	7	140348	14	348	858123567	0	!
	1 2						

The first and second tokenized messages (message IDs 267 and 68) each include one special parameter hence the s_parm_present parameter of those messages is set to one indicating the number of special parameters present. Specifically, those messages have a process ID parameter as a first parameter value. As noted above, the values of special parameters precede all other parameter values in the parameter portion of the message. Special parameter values are placed ahead of positional parameters to simplify parsing of the message template and generation and presentation of messages therefrom. In particular, the first two messages have a process ID value of P1237 as the special parameter value. Positional parameter values follow the special parameter value.

An English language message catalog may include the following entries related to the above tokenized messages:

```
267   "BEGIN %s %s (pid=%P) (jobid=%s)"
68    "%s started for user \"%s\".  (pid=%P)"
24    "Source:            %s:%s"
25    "Target:            %s:%s"
33    "Summary of Execution Phase:"
61    "        %s Configured         %s%s"
65    "        %s Skipped (in analysis)    %s%s"
347   "%d of %d filesets were Skipped."
348   "%d of %d filesets had no Errors or Warnings."
```

It will be noted that the first two of the above exemplary templates have a process ID special parameter format conversion specification embedded in the format string. Specifically, the "%P" element of those format strings is to be replaced by the special parameter process ID value from the tokenized message parameter portion. The purposes of this particular special parameter is, as noted above, to present the process ID of the process which generated the tokenized message entry (not otherwise readily available to the viewer/presentation program).

An equivalent message catalog for Japanese is depicted in FIG. 7. The Japanese exemplary catalog is displayed graphically in FIG. 7 due to limitations on printing of such Japanese symbols within this text body. Note that parameters in the templates for message ID 347 and 348 use positional designations to alter the English standard order of the parameters provided. The Yen symbol "¥" is used at the end of lines in FIG. 7 which are continued to the next line of the message template display (analogous to the English use of the backslash "\" character).

When presented on a system configured for English as the local native language and thereby selecting the proper message catalog including the above exemplary templates the above tokenized messages would be displayed or printed (presented) as follows:

```
=======   03/11/97 16:37:48 MST BEGIN install AGENT SESSION (pid=1237)
          (jobid=swposix-0312)
     *    Agent session started for user "root@swposix.fc.hp.com"
          (pid=1237)
     *    Source:     swi18n.fc.hp.com:/var/spool/sw
     *    Target:     swposix:/
     *    Summary of Execution Phase:
               Configured test_token.ok_filest,r=
               Skipped (in analysis) test_token.skip_filest,r=
     *    1 of 2 filesets were Skipped.
     *    1 of 2 filesets had no Errors or Warnings.
```

Certain information presented in the above exemplary English language presentation may be automatically generated by the viewing system in accordance with local standards based upon information encoded in the message header portion. For example, the "=======" string and the date and time displayed on the first message are generated by the viewer system based upon the message ID and message type fields of the message header portion of the tokenized message. The equal signs may be preferred locally to signify the start of a set of related logged messages. The corresponding date and time may be generated from the time stamp value in the message header. The "*" symbol is likewise generated by the viewer program locally in response to the message ID and message type fields of the header portion of the corresponding tokenized message entries. Those skilled in the art will recognize many such custom formatting options based upon information stored in the header portion of a tokenized message. Start and end of related logged messages, dates and time thereof, etc. all may be encoded in the message header values and may be used to generate particular formatted output corresponding to the tokenized message.

An equivalent Japanese presentation of the same tokenized log file messages is depicted in FIG. 8. The Japanese presentation is displayed graphically in FIG. 8 due to limitations on printing of such Japanese symbols within this text body. Note that positional parameters have been reversed relative to the English presentation for messages presented in accordance with templates for message IDs 347 and 348.

GENERATION OF TOKENIZED MESSAGES

In the simplest case, the tokenized messages may be generated by any of several standard programming techniques to simply write message header portion and parameter portion with the requisite values. As noted above, the values may be written as binary data or encoded as character data (e.g., ASCII encoded). The token values are localized as described above into appropriate formats for presentation in the locally preferred native language. However, in the preferred embodiment, a more rigorous method of generating the tokenized messages enables type checking and/or type conversion features in the generation of tokenized messages.

In the preferred embodiment, a function is invoked with parameters including the message ID of the message to be generated followed by a variable length (possibly null) list of parameter values to be substituted into the message template corresponding to the message ID. The message ID and parameter values are written to the tokenized message as noted above. However, before doing so, the methods of the preferred embodiment of the present invention retrieve the corresponding template from a canonical message catalog for purposes of generating any required special parameter values and for purposes of checking and/or converting the types of the variable length list of supplied parameters values.

FIG. 9 describes a preferred method of the present invention for generation of a tokenized message using templates in a canonical message catalog. The canonical message catalog is preferably the message catalog encoded in the language of the system or application designers. These message template are originally created to express a particular concept required by the system or application in the native language of the program designer. Once such an application or system is designed and operational, all such message template are translated into other required languages for support of multi-national deployment of the system or application.

Element 900 of FIG. 9 is first operable to initialize the header portion of the message to be generated. The header portion is initialized in a buffer used to construct the tokenized message requested. The structure of the header portion is as discussed above. Element 902 is next operable to locate and retrieve the message catalog entry corresponding tot he supplied message ID. The template of the retrieved entry is then parsed by element 904 to locate any special parameter format conversion specifications (as described above) and to generate appropriate special parameter values (if any) in the parameter portion of the tokenized message buffer under construction. For each special parameter conversion specification so located, the requisite parameter value is retrieved (or generated) and written to the parameter portion of the tokenized message being generated. If multiple special parameter format conversion specifications are so located, each parameter value (if required) is written to the parameter portion of the message in the order in which they are encountered.

Element 906 then updates the header portion of the tokenized message under construction to reflect the actual number of special parameter values present in the parameter portion of the message. Element 908 then appends the supplied parameter values (the positional parameters) to the end of the parameter portion of the tokenized message being constructed. The supplied values are written to the parameter portion in the order in which they are supplied to the function. Element 906 may also verify the types of the supplied parameter against the types identified in the template format specification corresponding to the positional parameters. The type of each supplied parameter is also supplied to the function or is otherwise determined in accordance with the standards of the programming language implementing the message generation function.

Element 910 is then operable to determine if the types of all parameters checked (if any) verify in comparison to the types expected by the retrieved message template. If all types verify or if type checking is not performed, element 912 is next operable to write the generated tokenized message buffer to the tokenized log message file (also supplied to the function of FIG. 9). In either case, element 914 then completes generation of the tokenized log file message entry and returns the completion status to the calling program.

While the invention has been illustrated and described in detail in the drawings and foregoing description, such illustration and description is to be considered as exemplary and not restrictive in character, it being understood that only the preferred embodiment and minor variants thereof have been shown and described and that all changes and modifications that come within the spirit of the invention are desired to be protected.

Claims

1. In a computing system (100-130, 200, 300), a method for manipulating messages in a language independent manner CHARACTERIZED IN THAT the method comprises the steps of:

   generating a first message in a language independent format (400-410, 900-914);
   selecting a message catalog corresponding to a preferred language (502); and
   translating said first message into a second message encoded in said preferred language in accordance with said message catalog (504-508, 600-608).

2. The method of claim 1 wherein the step of generating includes the step of:

   generating a message header portion of said first message (400, 900) wherein said message header includes a message ID identifying a corresponding entry in the selected message catalog.

3. The method of claim 2 wherein the step of translating includes the steps of:

   locating an entry in the selected message catalog corresponding to said message ID of said message header wherein said entry includes a template for translating said first message into said second message (502-506); and
   generating said second message in accordance with the template in the located entry of said selected message catalog (506, 600-608).

4. The method of claim 3 wherein the step of generating said first message further includes the step of:

   generating a parameter portion of said first message (402-408).

5. The method of claim 4 wherein said parameter portion includes at least one parameter value (404).

6. The method of claim 5 wherein the step of generating said second message includes the steps of:

   identifying in the template a variable portion thereof to be replaced by one of said at least one parameter value

of said parameter portion of said first message (600-602); and
substituting said one of said at least one parameter value for said variable portion of the template (604-606).

7. The method of claim 5 wherein each of said at least one parameter value is associated with unique positional indicia and wherein the step of generating said second message includes the steps of:

identifying in the template a variable portion thereof to be replaced by an identified one of said at least one parameter value of said parameter portion of said first message wherein said identified one of said at least one parameter value is identified by said unique positional indicia associated with said identified one of said at least one parameter value (600-602); and
substituting said identified one of said at least one parameter value for said variable portion of the template (604-606).

8. The method of claim 4 wherein the step of generating said second message includes the steps of:

identifying in the template a variable portion thereof to be replaced by a special parameter value (600-602);
deriving said special parameter value from information in said header portion of said first message (604); and
substituting said special parameter value for said variable portion of the template (606).

9. The method of claim 4 wherein said parameter portion includes at least one special parameter value and wherein the step of generating said second message includes the steps of:

identifying in the template a variable portion thereof to be replaced by one of said at least one special parameter value (600-602); and
substituting said one of said at least one special parameter value for said variable portion of the template (604-606).

10. The method of claim 4 wherein the step of generating said first message includes the steps of:

parsing said template to determine whether special format specifications for special parameters are present therein (904); and
generating special parameter values in said parameter portion of said first message in response to a determination that said special format specifications are present in said template (904-908).

11. The method of claim 10 wherein the step of generating said first message further includes the steps of:

parsing said template to determine whether standard format specifications for standard parameters are present therein (904); and
generating standard parameter values in said parameter portion of said first message in response to a determination that said standard format specifications are present in said template (904-908).

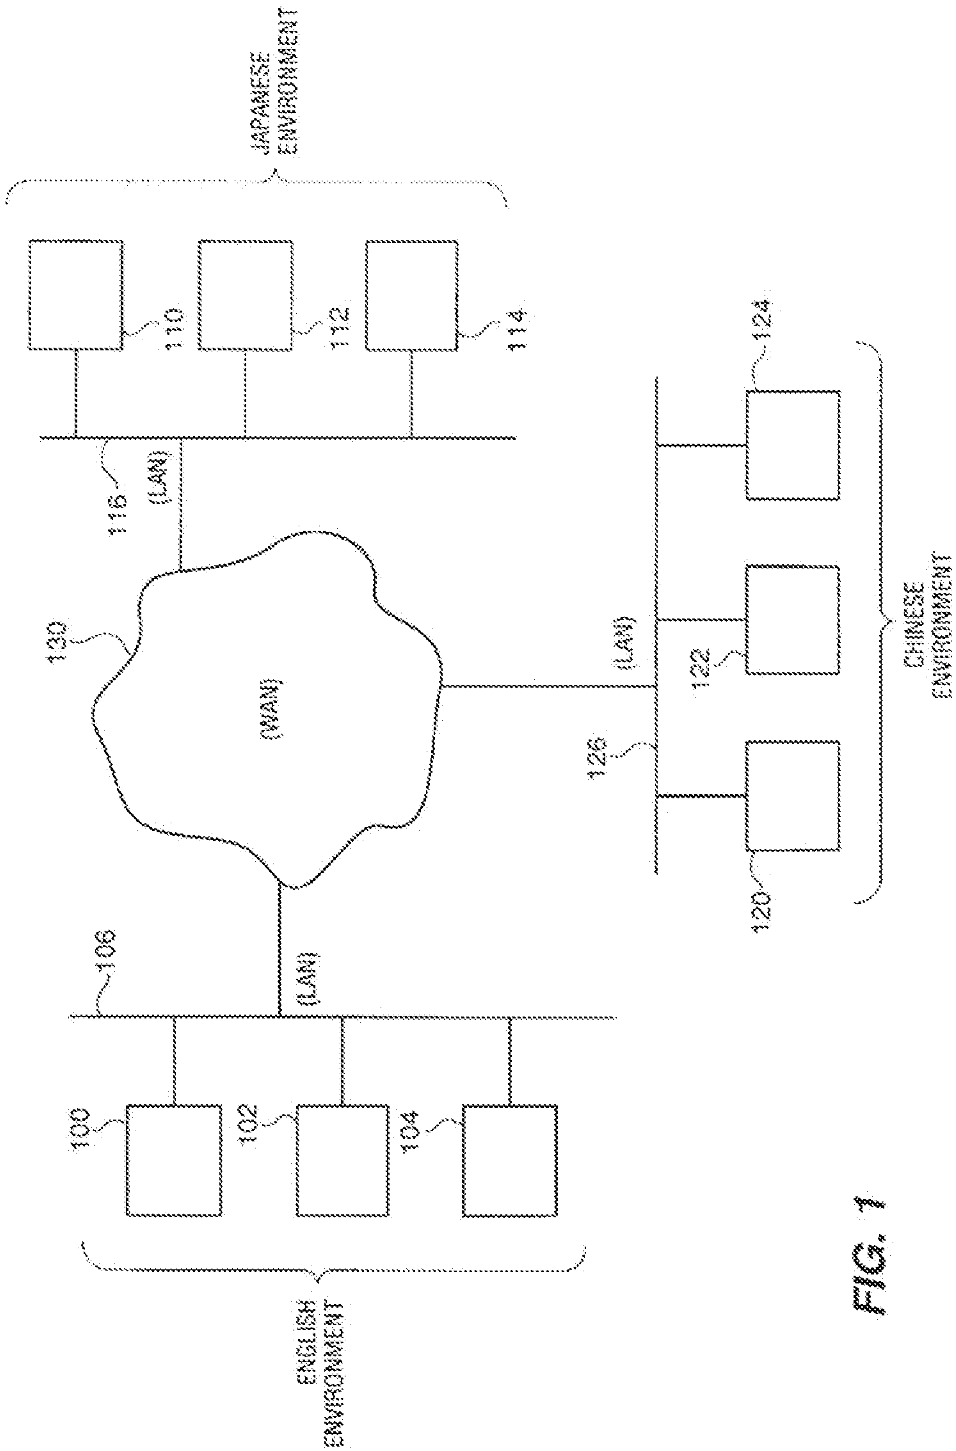
JAPANESE ENVIRONMENT
110
112
114
116
(LAN)
130
(WAN)
(LAN)
126
122
124
120
CHINESE ENVIRONMENT
106
(LAN)
100
102
104
ENGLISH ENVIRONMENT

FIG. 1

FIG. 2

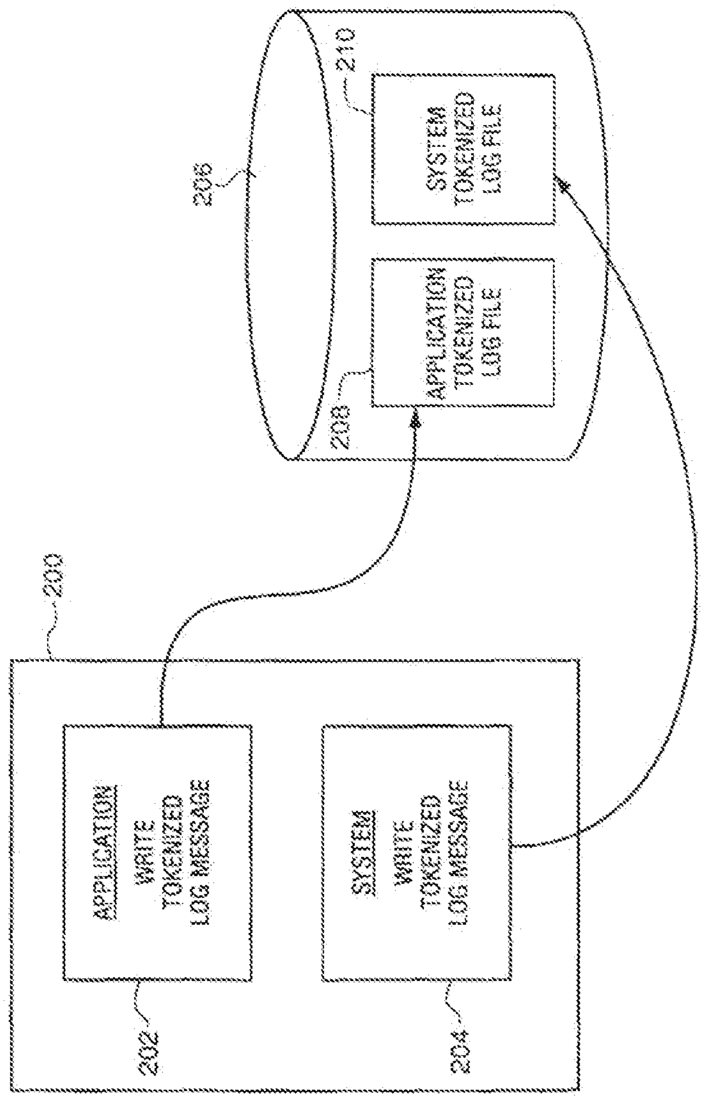
200
APPLICATION
WRITE
TOKENIZED
LOG MESSAGE
202
SYSTEM
WRITE
TOKENIZED
LOG MESSAGE
204
206
APPLICATION
TOKENIZED
LOG FILE
208
SYSTEM
TOKENIZED
LOG FILE
210

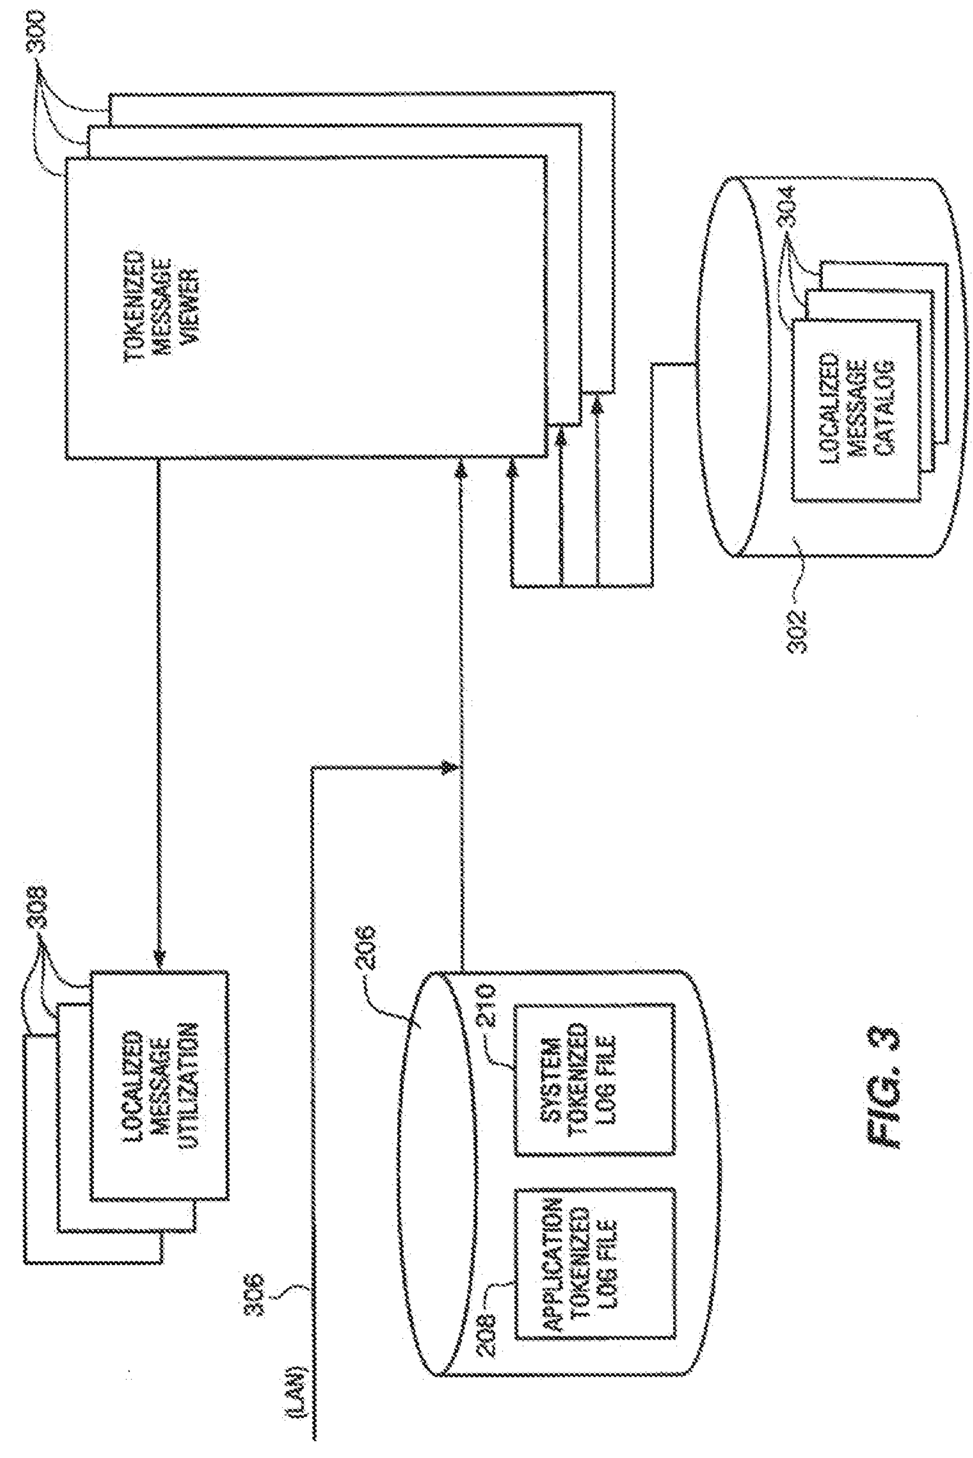
300
TOKENIZED MESSAGE VIEWER
304
LOCALIZED MESSAGE CATALOG
302
308
LOCALIZED MESSAGE UTILIZATION
206
210
SYSTEM TOKENIZED LOG FILE
208
APPLICATION TOKENIZED LOG FILE
306
(LAN)

FIG. 3

FIG. 4

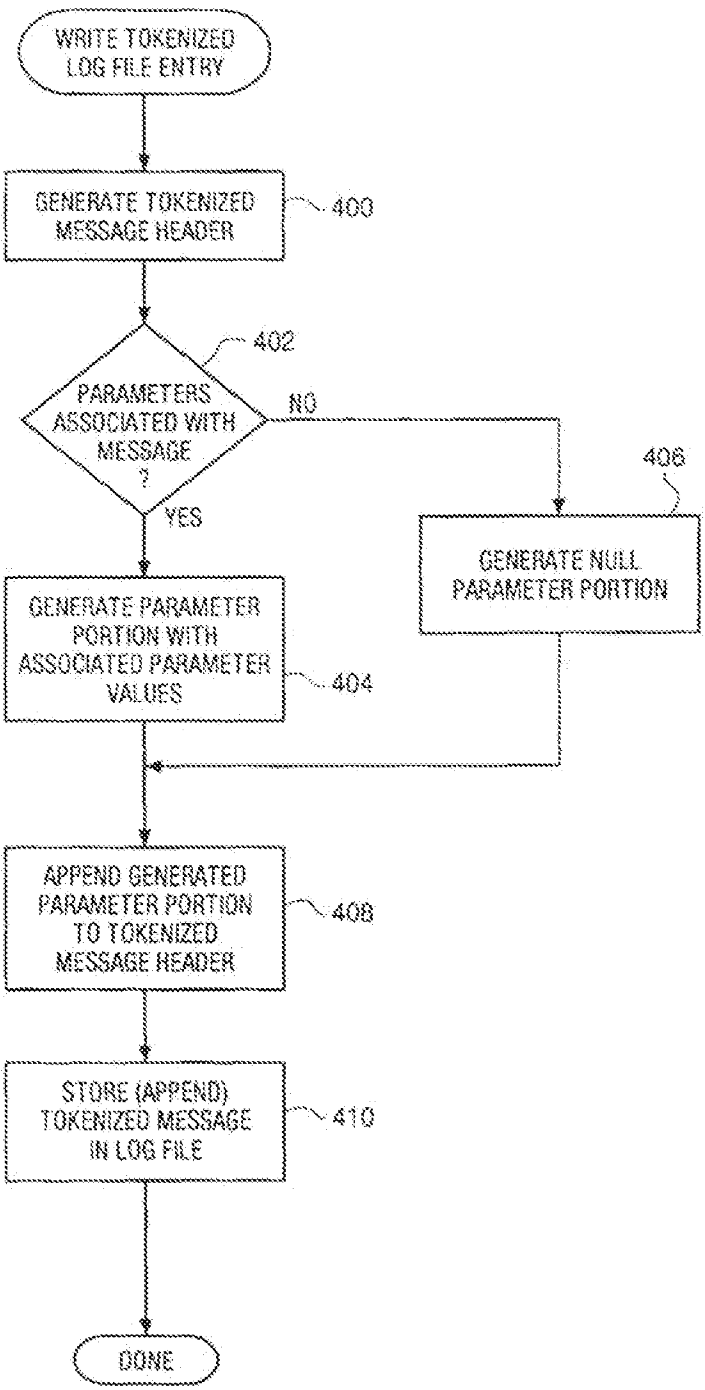
WRITE TOKENIZED LOG FILE ENTRY
GENERATE TOKENIZED MESSAGE HEADER
400
402
PARAMETERS ASSOCIATED WITH MESSAGE ?
NO
YES
406
GENERATE NULL PARAMETER PORTION
GENERATE PARAMETER PORTION WITH ASSOCIATED PARAMETER VALUES
404
APPEND GENERATED PARAMETER PORTION TO TOKENIZED MESSAGE HEADER
408
STORE (APPEND) TOKENIZED MESSAGE IN LOG FILE
410
DONE

FIG. 5

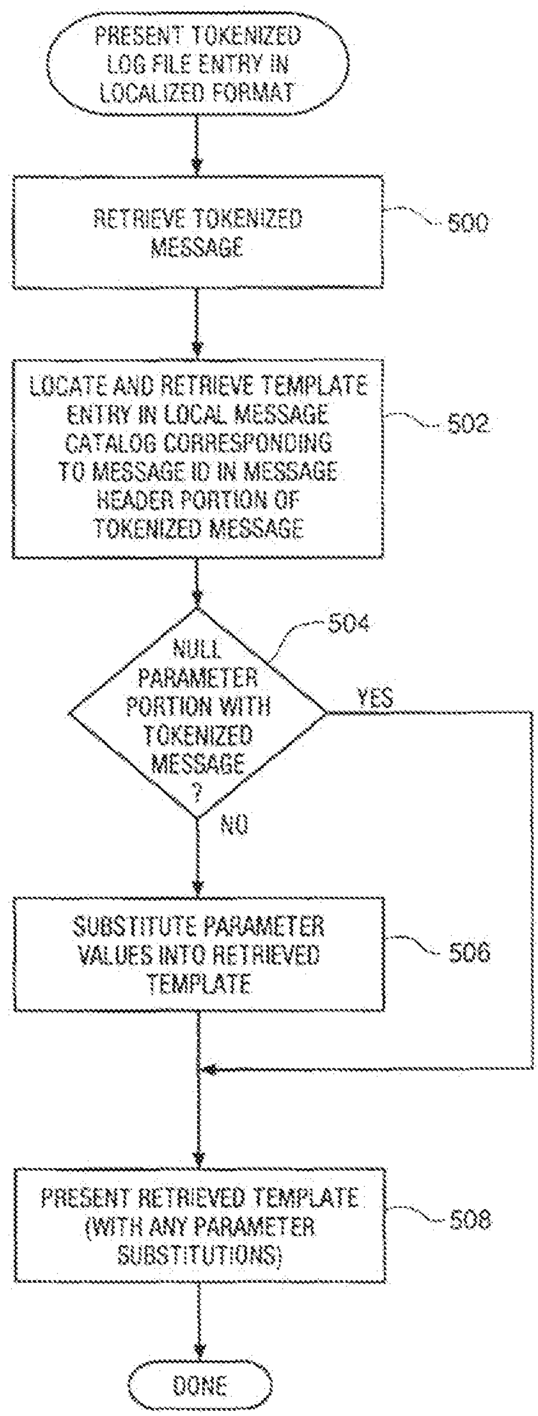
PRESENT TOKENIZED LOG FILE ENTRY IN LOCALIZED FORMAT
RETRIEVE TOKENIZED MESSAGE
500
LOCATE AND RETRIEVE TEMPLATE ENTRY IN LOCAL MESSAGE CATALOG CORRESPONDING TO MESSAGE ID IN MESSAGE HEADER PORTION OF TOKENIZED MESSAGE
502
504
NULL PARAMETER PORTION WITH TOKENIZED MESSAGE ?
YES
NO
SUBSTITUTE PARAMETER VALUES INTO RETRIEVED TEMPLATE
506
PRESENT RETRIEVED TEMPLATE (WITH ANY PARAMETER SUBSTITUTIONS)
508
DONE

FIG. 6

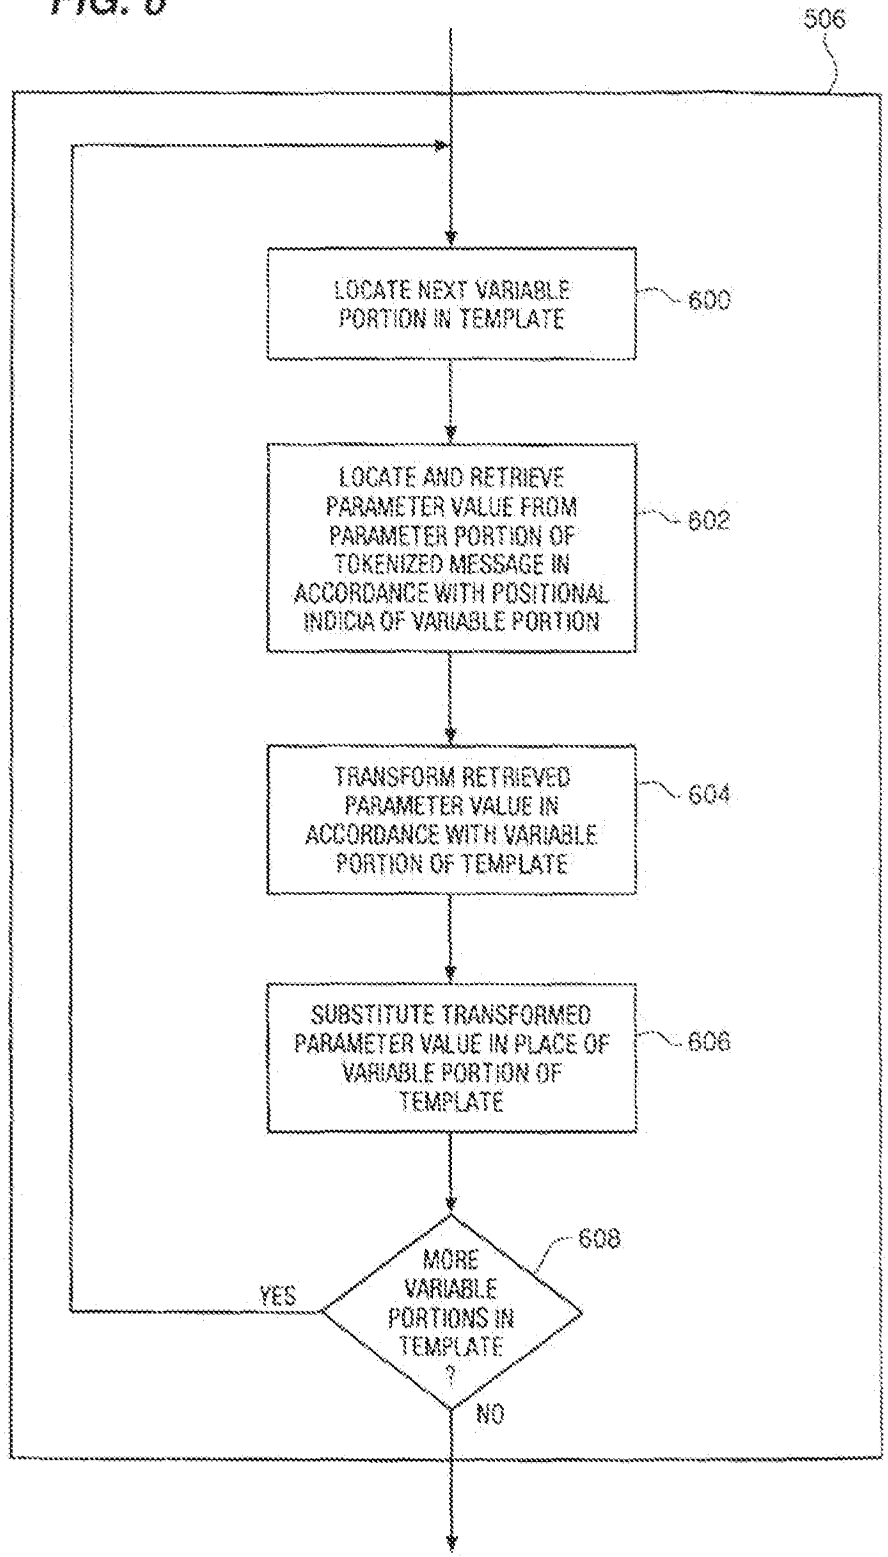
506
LOCATE NEXT VARIABLE PORTION IN TEMPLATE
600
LOCATE AND RETRIEVE PARAMETER VALUE FROM PARAMETER PORTION OF TOKENIZED MESSAGE IN ACCORDANCE WITH POSITIONAL INDICIA OF VARIABLE PORTION
602
TRANSFORM RETRIEVED PARAMETER VALUE IN ACCORDANCE WITH VARIABLE PORTION OF TEMPLATE
604
SUBSTITUTE TRANSFORMED PARAMETER VALUE IN PLACE OF VARIABLE PORTION OF TEMPLATE
606
MORE VARIABLE PORTIONS IN TEMPLATE ?
608
YES
NO

```
267 ¥
BEGIN %s %s (pid=%P) (jobid=%s)
68 %s を開始しました。 ¥n(ユーザー: ¥"%s¥" , pid=%P)
24 ソース:                      %s:%s
25 ターゲット:                  %s:%s
33 実行フェーズのサマリ:
61   %s 構成しました          %s%s
65   %s スキップしました (解析時に)¥n%s%s
347 ¥
%2$ld 個のファイルセットのうちの %1$ld 個がスキップされました。
348 ¥
%2$ld 個のファイルセットのうちの %1$ld 個に対してエラー、警告は¥n¥
発生しませんでした。
```

FIG. 7

```
=======  1997年03月11日 16時37分48秒 MST  BEGIN install AGENT SESSION
         (pid=1237) (jobid=swposix-0312)

       * Agent session を開始しました。
         (ユーザー: "root@swposix.fc.hp.com" , pid=1237)
       * ソース:                     swi18n.fc.hp.com:/var/spool/sw
       * ターゲット:                 swposix:/
       * 実行フェーズのサマリ:
             構成しました        test_token.ok_filest,r=
             スキップしました (解析時に)
         test_token.skip_filest,r=
       * 2 個のファイルセットのうちの 1 個がスキップされました。
       * 2 個のファイルセットのうちの 1 個に対してエラー、警告は
         発生しませんでした。
```

FIG. 8