

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
22 June 2006 (22.06.2006)

PCT

(10) International Publication Number
WO 2006/065503 A2

(51) International Patent Classification:
G06Q 10/00 (2006.01)

(21) International Application Number:

PCT/US2005/042866

(22) International Filing Date:

30 November 2005 (30.11.2005)

(25) Filing Language:

English

(26) Publication Language:

English

(30) Priority Data:

11/014,248 17 December 2004 (17.12.2004) US

(71) Applicant (*for all designated States except US*): **TACODA SYSTEMS, INC** [US/US]; 8th Floor, 345 Seventh Avenue, New York, NY 10001 (US).

(72) Inventors; and

(75) Inventors/Applicants (*for US only*): **BEYDA, Gil** [US/US]; 1445 Catlin Way, Dresher, PA 19025 (US). **BALASUBRAMANIAN, Krishna** [IN/US]; 1843 Acorn Lane, Abington, PA 19001 (US). **LOCKER, Donald** [US/US]; 14 Stoneleigh Drive, Lebanon, PA 17042 (US). **MCDONALD, Kevin** [US/US]; 32 Pagett Street, Charleston, SC 29492 (US). **WILSON, Joseph, G.** [US/US]; 9337 Chestnut Knolls Drive, Fairfax, VA 22302 (US). **MORGAN, David, R.** [US/US]; Apartment 1B, 55 West 74th Street, New York, NY 10023 (US). **LOWRY, Alison, P.** [US/US]; 242 Bryson Road, Atglen, PA 19310 (US).

(74) Agent: **YOHANNAN, David, R.**; Collier Shannon Scott, PLLC, Suite 400, 3050 K Street, Northwest, Washington, DC 20007-5108 (US).

(81) Designated States (*unless otherwise indicated, for every kind of national protection available*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, LY, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Published:

— *without international search report and to be republished upon receipt of that 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: AUDIENCE MATCHING NETWORK WITH PERFORMANCE FACTORING AND REVENUE ALLOCATION

(57) Abstract: Performance based delivery of content to an audience member. A network of audience member information collection domains provide information about audience members. A profiled audience member may belong to a network segment. Delivery of an advertisement to the profile audience member is accommodated based upon membership in the network segment and performance criteria. The performance criteria may be configured to benefit a publisher, such as by maximizing revenue. Allocating revenue based upon the collection of data used to target audience members is also provided.



WO 2006/065503 A2

**AUDIENCE MATCHING NETWORK
WITH PERFORMANCE FACTORING AND REVENUE ALLOCATION**

Inventors:

Gil Beyda
Krishna Balasubramanian
Donald J. Locker
Kevin M. McDonald
Joseph G. Wilson
David R. Morgan
Alison P. Lowery

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application is a continuation-in-part of U.S. Patent Application No. 10/981,733 filed on November 5, 2004, and entitled "Audience Targeting System with Segment Management," which is a continuation-in-part of U.S. Patent Application No. 10/669,791, filed on September 25, 2003, and entitled "System and Method for Segmenting and Targeting Audience Members," which claims the benefit under 35 USC §119 of Provisional Patent Application No. 60/491,521, filed on August 1, 2003. The entire contents of these Applications are hereby incorporated by reference.

BACKGROUND OF THE INVENTION

1. Field of the Invention

[0002] This invention relates generally to audience targeting and more particularly to matching an audience with deliverable content such as advertising.

2. Description of the Related Art

[0003] Targeted marketing has long been known as an effective method for reaching consumers. When the consumer receives only relevant content (advertisements, etc.) from a provider, the consumer is more likely to patronize the particular provider, make purchases, and provide additional personal information that may assist in refining the provider's "view" of the consumer. As such, targeted marketing can lead to a more focused and robust interaction with the consumer. This, correspondingly, can lead to a more rewarding interaction for the provider by generating increased revenue.

[0004] In order to effectively target a consumer, it may be desirable for marketing systems to react to consumer information received from a variety of online and offline sources. These sources may include databases and servers, as well as multiple web properties within a network of affiliated websites. Moreover, the consumer information may be collected from a variety of sources in diverse formats. It may also be desirable for marketing systems to interact with the systems that actually deliver the content to the user. In short, an effective marketing system may appreciate the characteristics and preferences of a specific user regardless of the number or type of channels through which contact with the user is made.

[0005] Some known systems, however, are only adapted to receive information from a single source (e.g., registration information provided by the consumer). Other systems may receive information from multiple sources, but are unable to usefully combine information relating to the same consumer and communicate it to the necessary content delivery system. Thus, it may be desirable to have a system and method for delivering content that integrates with and aggregates data from various sources, including the underlying systems that deliver content to the consumer.

[0006] Known systems for delivering targeted content to consumers are focused on reaching the greatest quantity of consumers, without considering the value of interacting with

each particular consumer. For example, some systems may deliver “targeted” content to each member of a group of consumers based on the fact that each subscribes to the same magazine. These systems, however, do not consider that only a portion of the group may make on-line purchases, for example, in addition to subscribing to the magazine. This failure to recognize and differentiate “valuable” consumers can lead to lost revenue for the content provider. In addition, the delivery of content to a significant volume of low-value consumers may expend valuable system resources. Accordingly, it may be desirable to have a means of delivering the appropriate content to the appropriate user in order to maximize the value of the relationship between the provider and the consumer.

[0007] Another problem with content delivery systems is that the information upon which targeting is based may rapidly become stale. An audience member deemed to have particular characteristics may no longer have such characteristics by the time content is delivered. New potential audience members may also become available after determination of a targeted group. The volatility of audience member characteristics and the high volume of information to be processed both present difficulties to systems that seek to target well tailored audiences. Content delivery systems are also often myopic, merely carrying out content delivery as dictated by the particular domain in which the system resides. This prevents appreciation of activities in other domains.

[0008] Still another problem with content delivery systems, particularly those that seek to collect information and deliver content to particular audience members over the Internet, is the potential for faulty identification of audience members. For example, some systems may use cookies to attempt to uniquely identify an audience member. This presents potential problems because a given person may use several computers and thereby generate several cookies. Software and browser maintenance activities may also prompt the deletion of cookies. These various factors may prompt the proliferation of unnecessary and sometimes erroneous profiles.

[0009] It is also difficult for publishers to serve advertisements such that revenue is maximized, or accommodate proper serving of advertisements by third party providers. Finally, the allocation of credit and corresponding revenue for activities related to the serving of advertisements remains inadequate.

SUMMARY OF THE INVENTION

[0010] The present invention accommodates the delivery of content such as advertisements to audience members.

[0011] In one aspect, performance based delivery of content to an audience member comprises managing audience segments corresponding to a network of audience member information collection domains, determining that a profiled audience member is a member of a network segment that includes audience members having a set of attributes corresponding to profile data collected in the network of audience member information collection domains, and accommodating the delivery of an advertisement to the profiled audience member based upon membership in the network segment and performance criteria.

[0012] The performance criteria may be configured to benefit a publisher, and may further seek to maximize revenue to the publisher. For example, accommodating the delivery of an advertisement may involve selecting an advertisement with the highest available payment to the publisher.

[0013] Placement of advertisements may further entail receiving from the publisher a set of delivery criteria corresponding to potentially deliverable advertisements, and selecting an advertisement with the highest available payment to the publisher based upon a comparison of current delivery metrics to the set of delivery criteria.

[0014] Another aspect optimizes the delivery of advertisement inventory served by a third party. Optimizing the delivery of advertisement inventory served by the third party may include

providing information about membership in the network segment in connection with an advertisement serving function of the third party.

[0015] Still another aspect of the present invention provides network segment definition through examination of the data profiles for audience members that have engaged in a desired behavior, such as responding to an advertisement.

[0016] Still another aspect of the present invention supports allocation of revenue corresponding to the delivery of content to an audience member, and may comprise managing audience segments corresponding to a network of audience member information collection domains, determining that a profiled audience member is a member of a network segment that includes audience members having a set of attributes corresponding to profile data collected in the network of audience member information collection domains, accommodating the delivery of an advertisement to the profiled audience member based upon membership in the network segment, and allocating revenue corresponding to the delivery of the advertisement based upon participation as a data provider in the collection of information.

[0017] The present invention can be embodied in various forms, including business processes, computer implemented methods, computer program products, computer systems and networks, user interfaces, application programming interfaces, and the like.

BRIEF DESCRIPTION OF THE DRAWINGS

[0018] These and other more detailed and specific features of the present invention are more fully disclosed in the following specification, reference being had to the accompanying drawings, in which:

[0019] FIG. 1 is a block diagram illustrating an example of a system for delivering content to an audience member.

[0020] FIG. 2 is a flow diagram illustrating an example of delivering content to an audience member.

[0021] FIG. 3 is a flow diagram illustrating an example of generating audience member profiles.

[0022] FIG. 4 is a flow diagram illustrating an example of tracking website pages visited by an audience member using a unique identifier.

[0023] FIG. 5 is a flow diagram illustrating an example of grouping audience members into segments for receipt of targeted content.

[0024] FIG. 6 is a flow diagram illustrating an example of directing targeted content to audience members in a segment.

[0025] FIG. 7 is a block diagram illustrating another example of a system for delivering content to an audience member.

[0026] FIG. 8 is a block diagram illustrating an example of an audience targeting system that includes segment management.

[0027] FIGs. 9A-B are respectively a block diagram illustrating an example of a particular extractor 900 and a schematic diagram that exemplifies a model for extracting profile data.

[0028] FIGs. 10A-B are schematic diagrams illustrating an example of a segment management architecture and corresponding calculation of segments.

[0029] FIGs. 11A-B are schematic diagram illustrating an example of processing data tables to manage and produce segments.

[0030] FIG. 12 is a block diagram illustrating an example of an audience targeting system that includes profile synchronization.

[0031] FIG. 13 is a flow diagram illustrating an example of a process for profile synchronization.

[0032] FIGs. 14A-B are schematic diagrams illustrating an example of a network for matching an audience with deliverable content according to the present invention.

[0033] FIG. 15 is a block diagram illustrating an example of an audience matching network system according to the present invention.

[0034] FIGs. 16A-B are event diagrams illustrating an example of a computer implemented process for matching audience members to deliverable content according to the present invention.

[0035] FIG. 17 is a block diagram illustrating an example of an audience matching network system that includes advertising revenue and performance management according to another aspect of the present invention.

[0036] FIG. 18 is a flow chart illustrating an example of delivering advertisements to with revenue and performance management according to this aspect of the present invention.

[0037] FIG. 19 is a block diagram illustrating an example of an audience matching network system that includes universal profile synchronization (UPS) according to another aspect of the present invention.

[0038] FIG. 20 is a flow diagram illustrating an example of a process for UPS.

DETAILED DESCRIPTION OF THE INVENTION

[0039] In the following description, for purposes of explanation, numerous details are set forth, such as flowcharts and system configurations, in order to provide an understanding of one or more embodiments of the present invention. However, it is and will be apparent to one skilled in the art that these specific details are not required in order to practice the present invention.

[0040] One embodiment of the system 10 for delivering content to an audience member is shown in Fig. 1. The system 10 includes a first server 120 which hosts an extractor program

122. The first server 120 is operatively connected to one or more offline databases 110, and one or more external content servers 160. The offline databases 110 and external content servers 160 are also operatively connected to one or more web servers 170. The web servers 170 may provide website pages to an audience member computer 180 in a conventional manner. The web servers 170 are also operatively connected to a targeting engine program 152 resident on a second server 150. The first and second servers 120 and 150 may be operatively connected to a third server 130 which contains a database 132 (referred to as the data warehouse) for storing audience member profile data. In some embodiments of the present invention, the same server may act as the first, second, and/or third servers 120, 150, and 130. A control console 140 may be operatively connected to the third server 130.

[0041] FIG. 2 is a flow diagram illustrating an example of delivering content to an audience member. This may include three primary stages: data collection and profile generation; audience segmentation and analysis; and interface to external systems. During data collection and profile generation, offline data sources 110 are searched to collect profile data relating to individuals (referred to as audience members). This profile data is stored in the data warehouse 132. During audience segmentation and analysis, the profile data for audience members is used to categorize the audience members into segments. For example, profile data may indicate that a particular audience member subscribes to Golf Magazine, and thus has some interest in golf. That audience member may then be included in a segment (i.e., group) of audience members that are also interested in golf. During the interface to external systems stage, a targeting engine 152 may use the inclusion of the audience member in a segment to direct targeted external content to the audience members in that segment. Continuing with the example posed above, audience members in the “golf” segment may have golf related content sent to them as a result.

[0042] With continued reference to Fig. 2, data collection and profile generation may involve the offline databases 110, the extractor program 122, and the data warehouse 132. Initial profile information about individual audience members may be collected from available databases, such as a registration database 112 and a subscription database 114 by the extractor 122. Registration and subscription databases 112 and 114 may include audience member profile data collected as a result of the audience member registering with, or subscribing to, any type of service, including but not limited to an Internet, magazine, newspaper, newsletter, cable, telephone, or wireless service, for example. These registration and subscription databases may include a wide variety of profile information such as name, gender, age, birth date, home and work addresses, telephone numbers, credit and charge card information, marital status, income level, number and age of children, employment history, hobbies, recent purchases, educational status, interests, preferences, and the like, for example.

[0043] The extractor 122 is a program that is used to parse and identify audience member profile data from within a set of data. The extractor 122 may be constructed using Java, Perl, C++, C#, SQL, or any other similar programming language. The extractor 122 may be resident on a server 120, or multiple servers. The extractor 122 may be governed by a set of extraction rules that determine the source(s) and format(s) of profile data that may be used to create a profile for an audience member, and the categories of profile data to be collected. The extraction rules may include a series of text rules (using matching utilities such as string matching or regular expressions) that are used to transform data in one form into a more standardized form while removing unneeded data. The extraction rules may include, for example, a statement such as “if string contains ‘A’ then output result ‘B’.”

[0044] The extractor 122 is operatively connected to a database 132 referred to as the data warehouse 132. The data warehouse 132 may be provided on a second server 130, and may be used to store the profile and segment affinity data relating to audience members. The extractor

122 may routinely update the profile and segment affinity data in the data warehouse 132. As new or modified profile data becomes available from the offline databases 110, the extractor 122 may modify the profile data for an audience member. The extractor 122 may also receive profile data directly from the audience member computer 180 and/or the targeting engine 152 that indicates the website pages visited, the web searches conducted, and the emails received by the audience member.

[0045] FIG. 3 is a flow diagram illustrating an example of generating audience member profiles. The steps shown in Fig. 3 show the manner in which the extractor 122 obtains profile data indicating the online website pages visited by an audience member. In step 210 the extractor searches the offline databases, such as registration and subscription databases, for profile data relating to individual audience members. The search of the offline databases may be initiated by an instruction received from the console 140. For example, an instruction could be given to collect profile data for all audience members who subscribe to the New York Times. Such an instruction necessitates that the extractor 122 have access to the subscription database for the New York Times.

[0046] The extraction rules determine the profile data that is collected. In step 212, the profile data extracted from the offline sources may be stored in the data warehouse. As there may be a need to determine the profile data that is associated with a particular audience member, the extractor may assign a unique identifier to the profile data in step 214. The unique identifier may be a string of numeric, alphabetic, alphanumeric, or other characters that may be used to identify one audience member.

[0047] In step 216, the unique identifier may be used to identify content visited by the audience member. The unique identifier may be so used by including it in a domain cookie associated with each website page visited by the audience member. Each of these domain cookies may be stored on the computer associated with the audience member, and may be used

to identify each particular website page visited by the audience member as being associated with the unique identifier. In step 218, the extractor may determine the domain cookies that are stored on the audience member's computer. Because these domain cookies include the unique identifier that identifies the particular audience member, the extractor may use these cookies to modify the profile data for a particular audience member to reflect that the audience member visited the website pages associated with the cookies. By combining the profile data obtained from the offline databases with the profile data updates that occur as a result of the audience member visiting website pages, a complete set of profile data may be collected for an audience member, reflecting both offline and online behavior and characteristics for the audience member.

[0048] Tracking the online history of an audience member requires that the system be able to uniquely identify audience members. This tracking may be accomplished by combining a unique identifier for each audience member with website pages in the network that the audience member has visited.

[0049] A method of providing the unique identifier in each of the domain cookies associated with a number of related website pages is illustrated in FIG. 4. Each of the domain cookies associated with the website pages visited by the audience member may be modified to include the unique identifier by designating one of the related website page domains as the primary website domain. A primary domain cookie with the unique identifier is established for the primary website domain. Usually, a network will already have a domain that can be used for this purpose. If not, one of the domains in the network may be designated as the primary domain.

[0050] With reference to FIG. 4, an audience member browser 300 initiates the process in step 340 by requesting a website page from a site within the network, www.domain1.com 310. Responsive to the website page request directed to www.domain1.com 310, a page is returned to

the browser 300 with an image tag which may reference the targeting engine 152 at te.domain1.com in step 342. In step 344, an image request is sent from the browser 300 to the targeting engine 152. If a unique identifier is not included in the request, in step 346 a redirect is sent to the browser 300 to the targeting engine 152 now referenced as te.primarydomain.com. The redirect includes a reference to the original targeting engine reference in step 344, te.domain1.com. For example, the redirect may be http://te.primarydomain.com/blank.gif?te.domain1.com. In step 348, the browser 300 may send this redirect request to te.primarydomain.com. Responsive to this request, in step 350 a primarydomain.com cookie containing a unique identifier for the audience member is assigned to the browser 300. In step 352, a second redirect is made of the browser 300 to te.domain1.com, that may include the same unique identifier as set in the primary domain cookie. For example, the redirect may be http://te.domain1.com/blank.gif?tid=7dha6wlk9927sha. In step 354, the redirect request is returned with the originally requested image and a domain1.com cookie with the same unique identifier as the primarydomain.com cookie.

[0051] After the process illustrated in FIG. 4 is completed, an audience member visit to another website in the network, such as www.domain2.com, may result in a request for an image at te.domain2.com. If the Targeting Engine 152 does not detect a domain2.com cookie with a unique identifier following the image request, the Targeting Engine 152 may redirect a request to primarydomain.com for a cookie. Responsive to this request to primarydomain.com, the primarydomain.com cookie is read and a redirect is sent back to the browser 300 containing the unique identifier contained in the primary domain.com cookie. The unique identifier in the primarydomain.com cookie is the same as previously set. The requested image may then be sent to the browser 300 along with the domain2.com cookie which may have the same unique identifier as the primarydomain.com cookie. This process of providing a domain cookie with

the unique identifier is carried out each time the audience member visits a new website page for the first time so long as the new website is related to the other websites in the network from the viewpoint of the Targeting Engine.

[0052] The Targeting Engine 152 may be a standalone web server, running on Apache, and using a MySQL database on a shared server, although the Targeting Engine 152 may be variously realized using alternative software and separate servers for Apache and the database. The Targeting Engine 152 may direct the setting of an additional cookie that may contain one or more segment identifiers. These cookies may then be used by other servers, such as, for example, an ad server, an email server, a streaming media server, and/or a web content server, to deliver targeted content to a particular audience member based upon one or more segments in the cookie.

[0053] With renewed reference to FIG. 2, the audience segmentation and analysis stage may be carried out by the data warehouse 132. The data warehouse 132 may assign a particular audience member to one or more segments based upon common profile characteristics. A segment of audience members may be defined as a group of audience members to which the system user (such as an advertiser) desires to send the same content. For example, returning to the example discussed above, a segment of audience members may be defined as all audience members that the system user selects to receive a particular golf advertisement. The selection of the audience members for receipt of this advertisement may be based on one or more audience member characteristics in the profile data.

[0054] A method of associating an audience member with a segment is illustrated in FIG. 5. In step 220, the profile data attribute values of audience members who will qualify for inclusion in the segment may be defined by a set of segment rules. The segment rules may be selected using the console 140. Any of number and/or range of profile data attribute values may be used to govern qualification for a segment. In step 222, the data warehouse 132 may search

the profile data to determine the audience members that qualify for the audience segment. This search may be carried out at the request of the system user, and if desired, on a routine basis, such as daily. In this manner, membership in the audience segment may be maintained up to date. In step 224, the data warehouse 132 may store segment affinity data to indicate the audience members that are included in a particular segment. It is appreciated that the segment affinity data may indicate that an audience member is in more than one segment. The segment affinity data is defined by a set of rules based upon the behavior and characteristics in the audience profile. Once a set of rules that define the segment affinity data are identified, a segment identifier is assigned to that particular set of rules. This segment identifier is then sent to the Targeting Engine 152, along with the audience unique identifier assigned previously by the Targeting Engine 152. In step 226, when the Targeting Engine 152 is notified that an audience member has requested a website page in the network, the Targeting Engine stores a segment-targeting cookie on the audience member's computer. The segment-targeting cookie includes the segment identifier that identifies the segments that the audience member is included in. The method of storing the segment-targeting cookie on an audience member computer is described in further detail below in connection with FIG. 6.

[0055] Profile data for audience members may also be manually analyzed to build segments. With renewed reference to FIG. 2, the server or servers that host the Targeting Engine 152 and the data warehouse 132 may be operatively connected to the console 140. The console 140 may be used to designate the offline databases used to initially populate the data warehouse with profile information, to set the rules for collecting profile information, and to create and view reports showing audience member profile data, audience member segment affinity data, and audience member Internet activity.

[0056] A method of delivering targeted content to an audience member based on the segment affinity data is illustrated in FIG. 6. With reference to FIG. 6, an audience member

requests a website page in the network of related websites in step 230. The Targeting Engine is notified of the website page request in step 232. Responsive to the audience members request for a website page, in step 234 the Targeting Engine determines whether or not a domain cookie, associated with the requested website page, includes a unique identifier for the audience member. If a unique identifier is not identified, the Targeting Engine will provide a website domain cookie with a unique identifier as described above in connection with Fig. 4. Once a website domain cookie is provided with a unique identifier, in step 236 the Targeting Engine may determine whether or not a segment-targeting cookie is already associated with the audience member in the data warehouse. The segment-targeting cookie may include a segment identifier that indicates the segment(s) to which the audience member belongs. If segment affinity data is stored in the data warehouse for the audience member, then a segment-targeting cookie is created and stored in the audience member computer with the appropriate segment identifier in step 238. In step 240, content may be delivered to the audience member based on the segment identifier in the segment-targeting cookie stored in the audience member computer.

[0057] If no segment-targeting cookie is identified in step 236, the Targeting Engine may query the data warehouse for any segment affinity data associated with the audience member. If no segment affinity data is stored for the audience member, a default segment-targeting cookie may be stored in the audience member computer. The default segment-targeting cookie may automatically expire after some fixed period of time, such as one day for example.

[0058] Once a segment-targeting cookie is stored on the audience member computer, the Targeting Engine may periodically update it with new segment affinity data for the audience member. Updating may occur automatically at fixed intervals, and/or in response to modifications to the profile data for the audience member.

[0059] A wide variety of content may be provided to the audience member as a result of the segment-targeting cookie being stored on the audience member computer. With renewed

reference to FIG. 2, content may include, but is not limited to website page advertisements, pop-up advertisements, emails, or the like.

[0060] The system 10 of the present invention is adapted to segment and target audience members for delivering content to an audience member across a plurality of digital mediums. The digital mediums may be heterogeneous, and may include, but are not limited to, a website network, a cable system, a non-web based internet network, a wireless communications system, such as a cellular phone or RF network, and/or any digital medium in which the means for interfacing the audience member with the digital content is uniquely addressable. It is contemplated that the digital medium may include other consumer technologies not yet developed.

[0061] FIG. 7 is a block diagram illustrating another example of a system for delivering content to an audience member. The system includes a digital cable network 400. The digital cable network 400 may include a home television having a uniquely addressable cable set-top box 480 as a means for interfacing the audience member with digital content. The digital cable network 400 may further include a cable head-end 450 for delivering segment targeted content to the set-top box 480. As will be apparent to those of ordinary skill in the art, the head-end 450 may include means for receiving a digital signal, such as, for example, a satellite receiving antennae, from a programming processor 460. The programming processor 460 programs the content to be delivered to the audience member, and provides the appropriate digital signal to the head-end 450. The programming processor 460 may be in communication with a cable company database 430 which may store, for example, subscription data relating to the audience member. The data may include a unique identifier of the audience member within the cable network 400. The programming processor 460 may interface with the system 10 of the present invention through a cable network/Internet bridge 440. As discussed above, the system 10 may include an audience member profile.

[0062] The digital cable network 400 may further include a cable company website provided by a web server 470 and accessible by the audience member via the Internet. The audience member may access the website 470 to request a service, such as, for example, ordering a movie, placing a repair order, and changing the level of cable service. The audience member may access the website 470 by providing the audience member's cable network identifier.

[0063] The system of FIG. 7 may be operated as follows for delivering content to an audience member across a plurality of digital mediums. The audience member may visit a website provided by a web server 170. The web server 170 may receive a request for content from the audience member, and provide website pages to an audience member computer 180 in a conventional manner. The website 170 may be owned by, or affiliated with, the owner of the cable network 400 and the website 470. The audience member may visit other sites related to the website 170 within a network. If necessary, a unique audience member identifier related to the website network is assigned to the audience member, and profile data is collected and stored, substantially as described above in connection with Figs. 3 and 4. The audience member may be associated with an audience segment defined by a set of segment rules substantially as described above in connection with FIG. 5.

[0064] The audience member may visit the website 470 to request a service from the cable company, at the same time providing the audience member's unique identifier within the cable network 400. The programming processor 460 may read the audience member's web network identifier, and associate the audience member's cable network identifier with this identifier. The programming processor 460 may then access the system 10 through the bridge 440, and accesses the segment affinity data relating to the particular audience member using the web network identifier. Based on the audience segment affinity data, the programming processor 460 defines the programming rules for the audience segment within the cable network 400. The appropriate

digital signal is then sent to the cable head-end 450, and the head-end 450 delivers the audience member targeted content via the set-top box 480 and the audience member's home television. The preferences and behavior of the audience member within the network 400 may also be used to update the member's profile within the system 10. In this manner, the audience member's preference and behavioral data is synchronized across a plurality of mediums into a common profile, and the content delivered to the audience member via those mediums may be customized based upon the characteristics of the profile.

[0065] FIG. 8 is a block diagram illustrating an embodiment of an audience targeting system 800 that includes a targeting engine (TE) 810, extractor (Extractor) 820, segment manager (SM) 830, and data warehouse 850.

[0066] The audience targeting system 800 and its components are illustrated collectively for ease of discussion. As described previously, the various components and corresponding functionality may be provided individually and separately if desired, such as by different servers that are assigned to the functionality of one or more of the components.

[0067] The functionality of the audience targeting system 800 is preferably provided by software that may be executed on any conventional processing system, such as those previously named or others. In that regard, the audience targeting system 800 may in turn be a component of a computer system containing a processor and memory. Although one modular breakdown is shown, it should be understood that the described functionality may be provided by greater, fewer and/or differently named components. Although a software embodiment is described, the audience targeting system 800 may also be provided as hardware or firmware, or any combination of software, hardware, and/or firmware.

[0068] As previously described, audience segments may be variously calculated, such as on a periodic basis. One model for accommodating audience segment calculation is a batch processing model. For example, at 24 hour intervals the Audience Targeting System 800 may

prompt a recalculation of all necessary audience segments based upon previously extracted data as well as any newly extracted data that had been discovered since the previous batch process. While this model is useful for many applications and for certain types of extractable data (e.g., data from registration sources, surveys and 3rd party data), it is not always the best model to implement. One issue with the batch processing model is that it can become computationally expensive, particularly where audience segments are recalculated based upon not only previously extracted data, but the newly extracted data. Another issue is that certain data sources may contain data that should be acted on more frequently than dictated by the batch processing interval. The example of the Audience Targeting System 800 illustrated in FIG. 8 accommodates what is referred to as a continuous processing model, although the batch processing alternative is also applicable to embodiments of the present invention.

[0069] The regular mining and sending of data to the Extractor 820 may be supported by what is referred to as “dock and shuttle” data extraction described further in connection with FIG. 9 below. The segment manager 830 and corresponding architecture is described further in connection with FIGS. 10A-B below. Recalculation of audience segments based upon incremental data, and processing data tables to manage and produce audience segments are described further in connection with FIGS. 11A-B. Each of these features may be provided in conjunction with the audience matching network aspects of the present invention that are described in connection with FIGS. 14A-B through 17 below.

[0070] Still referring to FIG. 8, the Audience Targeting System 800 accommodates the collection and coordination of data across multiple sites, as well as the targeting of audience members. In that regard, a user that wants to target a particular audience defines audience segments of interest. The audience segments correlate to user profile data that may comprise both characteristic and behavioral data. The characteristic data is often found in registration data and includes attributes such as age, gender, ZIP code, and household income. On the other

hand, behaviors may include attributes such as which sections were viewed on a site (e.g., sports, entertainment, health), which advertisements were seen (e.g., mortgage rates, allergy medication), referrers (e.g., AOL, Yahoo), the time of visiting the site (point in time, or range), and the frequency of visits to the site. Audience segments may be defined based upon such user profile data. In turn, the audience segments form the basis for the information that is extracted for analysis, reporting and targeting audience members in relevant segments.

[0071] Audience targeting is not limited to web applications. For example profile data might include behavioral attributes such as programs viewed, time viewed, etc., and characteristic attributes such as subscriber IDs or the like in applications involving a television set top box.

[0072] The TE 810 provides the means for assigning and coordinating unique identifiers corresponding to individual audience members. As previously described, when an audience member logs onto a page for the first time, the TE 810 places a cookie on their browser, which contains a unique identifier. Whenever that audience member returns to the site, the unique identifier is sent back to the TE 810. Based upon the unique identifier, the Audience Targeting System can set a segment cookie, which can be used for the delivery of targeted content such as ads, e-mails, etc. to the audience members computer or other relevant device. The TE 810 may also create logs of this activity. The unique identifier may be referred to as a profile identifier (PRID).

[0073] Another example of an extractor 900 is further described with reference to FIGs. 9A-B, which respectively are a block diagram illustrating an embodiment of an extractor 900 and a schematic diagram that exemplifies a model for extracting profile data. Although particular terms such as dock and shuttle are used because they are helpful in conceptually illustrating this aspect, it is noted that various alternative terminology may be used for elements that perform the same functions.

[0074] The Extractor 900 includes a shuttle 902, dock 904 and extraction module 906. The functionality of the so-configured Extractor 900 is best understood with concurrent reference to FIG. 9B, which also refers to other elements. The shuttle 902 may be code that resides on the data source. Its purpose is to mine local data locally and send it to the extractor (more specifically, the dock 904 on the extractor). In one embodiment, the shuttle 902 accomplishes this by assembling boxes. The dock 904 receives boxes and, when sufficient boxes are available (or a timeout occurs) creates a pallet 908 out of the boxes. The extractor works on those pallets 908. In order to accomplish this reliably, it is useful for the shuttle 902 to know where the source data resides. For the previously mentioned batch processing embodiments, it is also useful for the shuttle to handle the situation where log files “roll” and are archived by the customer. In this regard, the shuttle 902 interfaces with log data such as that provided by conventional log file generating elements (e.g., Apache).

[0075] A data agent may also be employed to assist in the gathering of information from website visitors. This may be provided in the form of code that is added to those pages in connection with which data collection is sought. The code may have header and function call portions that respectively identify the functions and variables that it needs to operate and ensure that all variables have been collected. The data agent may be configured to produce log lines suitable for receipt and processing by the TE. Examples of parameters include the version of the data agent, the page referrer, the page URL, time information, and the PRID. As will be described below in connection with profile synchronization, a REGID parameter may be provided as well. In addition to association with PRID as described, a cookie may delineate a unique REGID for an audience member in the same fashion. Another “cookie list” (CLIST) parameter may be used to identify the list of cookies that should be captured.

[0076] The dock 904 is the receiving area on the Extractor that manages the ordering and processing of pallets. Data from the shuttle 902 may be grouped into what is referred to as

boxes. Generally, a box contains a single event, but in some cases (e.g. OAS logs) a single record may contain several events. An event may be a time tagged user action on a source server. Examples of events may include a web page view, an ad impression, etc. A pallet 908 may be a collection of boxes, and is typically a collection of data mined from the data source and packaged for delivery to the extractor dock 906.

[0077] Various data sources may be supported by this model, but in one embodiment web log data is the data source. The shuttle 902 may be a persistent C++ application that processes data from a log file or pipe. Upon startup, the shuttle 902 finds the current log file (or pipe) and opens it for reading. In addition, the shuttle 902 establishes a connection to the dock 904 in order to be able to deliver pallets 908 to the extraction module 906 for processing.

[0078] The shuttle 902 may be configured to process data in a persistent loop until an unrecoverable error or external termination signal occurs. During the processing loop, the shuttle 902 reads up to a configurable number of available items (log lines) from the source and packages them into a box. If there are more items available than the maximum number of items, or if the total size of the items are greater than the maximum box size, the extra lines are written into an overflow buffer and will be inserted first into the next box created.

[0079] Once the box has been created, the shuttle 902 sends the box to the dock 904, along with an indication of the size of the box for validation purposes. The extraction module 906 acknowledges and validates the box and responds with an acceptance signal before the shuttle 902 will drop the existing box and repeat the processing loop.

[0080] More than one shuttle 902 can connect to a given dock to allow for multiple machines which all serve the same data source (e.g., multiple web servers responding to a single domain via a load balancer). Data from different shuttles 902 in a given dock is sorted into bays. These bays contain the unprocessed data for a given data source from a given shuttle.

[0081] The extraction module 908 is preferably configured to handle each data source type, and may include sub-modules for each different data source type (*e.g.*, one for each of OAS, W3C, IIS, etc.).

[0082] Finally, the extraction module 906 is responsible for processing data as pallets from the dock 904 and creating the output that gets sent to the data warehouse 850 for final import processing. Basically, the extraction module 908 component performs extraction as described in connection with the previously described embodiment of the Extractor (from FIGs. 1-7). The processed data may be referred to as profile data. In one embodiment, the profile data may be organized and thus provided as fact tables that are described further below.

[0083] The segment management aspect is now further described with reference to FIG. 8, which illustrates the segment manager 830 to include a segment organization module 832 that includes a console management module 834, a segment generation module 836 that includes a new segment calculation module 838 and a segment recalculation module 840, and a reporting module 842.

[0084] The segment manager 830 accommodates the definition and management of segments corresponding to audience members based upon characteristic and behavioral information. The segments are organized according to a hierarchical logical tree based architecture that allows scalable segment management and accommodates incremental recalculation of segments.

[0085] The segment organization module 832 facilitates user-definition of audience segments according to this architecture. It operates in conjunction with the console manager 834 which provide interfaces that allow users to define and configure segments according to the same logical architecture. These interfaces may be in the form of panels that illustrate segments and combinations of segments to produce new segments which will be further understood upon explanation of the architecture below.

[0086] The segment generation module 836 generates segments comprising appropriate audience members based upon the so-defined audience segments. The new segment calculation module 838 calculates new segments, and the segment recalculation module 840 calculates existing segments, in particular taking incremental data and recalculating such segments, thus avoiding the need to fully calculate the segment as though it were new each time new data arrives.

[0087] The segment generation module 836 may be configured to process segments continuously (*e.g.*, as a Windows service). For each pass, the segment generation module 836 reads a table in the database warehouse 850 that catalogs segments, to determine which segments it should process on that pass. A type identifier associated with the segments may indicate whether the segments are to be calculated anew, and thus passed to the segment calculation module 838, or incremental, and thus passed to the segment recalculation module 840.

[0088] Finally, the reporting module 842 communicates with the segment organization 832 and segment generation module 836 and produces customizable reports. The designer is free to structure the reporting options as desired. One example of a report is a "Known Audience Inside/Outside" report, which reports on the behavior of an audience segment in the sections outside the section behavior that defines the segment. For example, An Inside/Outside report on viewers of the News section would show the audience members behavior inside news and compare that to all other sections of the site. This may be used to target valuable behavior on other parts of the site. Another example of a report is a "Reach and Frequency Report", which reports on the reach (total audience) and frequency (number of times seen) for one or more ad campaigns. The reporting module 842 may implement conventional reporting tools including but not limited to Crystal Reports as provided by Business Objects SA, San Jose California.

[0089] FIGs. 10A-B are schematic diagrams illustrating an example of a segment management architecture 1000(a-b) and corresponding calculation of segments according to another aspect of the present invention. As introduced above, the profile data includes attributes that are correlated to audience members, and is the basis of the audience segment definitions that are used to target audience members with advertisements and/or other content.

[0090] Profile data may also be organized as “facts” that have one or more attributes. For example an “Age” fact may have one attribute – Age. However, an “ID” fact may have several attributes such as the PRID or a registration identifier (REGID) that uniquely identifies registration at the site. A “Section” fact may contain attributes for the Section, top level Section (that is, if Section is /News/International/Politics, Top Level Section would be /News), second level section (/News/International), site (site that section belongs to) and full path (Site + Section).

[0091] Profile data and the individual attributes comprising the profile data may be categorized as being (1) Characteristics (e.g., Age, Gender, Household Income); (2) Behaviors (e.g., Page Views, Ad Clicks); (3) PRID; or (4) Business Unit ID, which describes the site that a behavior occurred on.

[0092] The attributes may also be said to have dimensions or values that may be defined in tables for ease of computation. Moreover, attributes may be further defined based upon whether they are single or multi-valued. For example, Age, Gender, HHI are characteristics for which an audience member will only have a single value (e.g., an audience member cannot be both Male and Female). Conversely, behaviors have multiple values per audience member and some characteristics (e.g., e-mail newsletters subscriptions) also have multiple values.

[0093] The hierarchical architecture facilitates efficient calculation of the membership of audience segments. Lists of audience members belonging to particular segments may be

maintained. These membership lists may be logically combined to determine the membership of dependent (*e.g.*, child) audience segments.

[0094] As indicated, the segment management architecture 1000a includes a series of attribute segments, namely Section 1002, Gender 1004, and Household Income (HHI) 1006 as provided in this example. Base segments have attributes with particular values that correlate to relevant attribute segments 1002-6. Base segments for any number of attributes could be provided (*e.g.*, different behaviors different sections; different gender, different HHI). The illustrated segments are “Visited News” 1010, “Male” 1012, and “HHI > \$100K” 1014. Each of these may be considered as separate and distinct segments. However, these segments may also be logically combined to create new segments that depend from them. For example, the segment “Males who have Visited News” 1020 comprises a logical combination of the Males 1012 and Visited News Last 1010 segments. Still further, a third level in the hierarchy of segments may be defined as “Males who have Visited News with HHI > \$100K” 1030, which comprises a logical combination of the previously described segment 1020 with base segment 1014 (HHI > \$100K). In this fashion, the system may variously organize segments, and this same organization can be used as the basis for guiding the user through the definition of segments via the console manager 834. Notably, there may be instances where a user defines a complex segment directly, wherein the system automatically generates the base and any intervening segments accordingly, to facilitate calculation and recalculation of segments.

[0095] For ease of illustration, a logical “AND” operation has been described, which basically provides the intersection of two parent segments. The segment manager 830 supports various additional logical operations or set expressions, including “EXISTS”, which inserts entries from one parent; “OR”, which inserts entries from the union of two parents; as well as “exclusive AND”, and “exclusive OR”. Attribute expressions may also be used, such as one which inserts entries from a given parent segment that match specified criteria.

[0096] In addition to providing improved organization of segments, the segment management architecture 1000a facilitates proper maintenance of a segment population where incremental profile data is processed, without requiring a full calculation of the segment. That is, introduction of the new information to the existing segment is accommodated through limited processing involving the new information, in lieu of calculating the segment based upon application of its definition to the cumulative set of data. To accommodate this, entry and exit rules are implemented. An “entry” corresponds to an introduction of audience members to a particular segment based upon the incremental data, and an “exit” corresponds to a removal of audience members from a segment. Entries are basically audience members found to currently meet the criteria, but whom are not yet associated with the previously calculated segment. Exits are the opposite – they are audience members found to no longer meet the criteria.

[0097] FIG. 10B illustrates an entry and exit 1032 functionality for the segment management architecture 1000b. As described above, the Extractor continuously populates the data warehouse with profile data that identifies various attributes. As indicated, a Gender’ attribute segment 1004’ is generated responsive to incremental profile data. This generally represents audience members that have attributes defined under the attribute segment “Gender” within the incremental profile data. Among those are the previously described “Male” segment 1012. In that regard, exit and entry membership lists are built. Specifically, all audience members identified as being male in the incremental profile data are provided in an entry membership list for the Male segment 1012. Similarly, all those audience members who do not have the relevant attribute (which may be referred to as “not male”) are provided in an exit membership list for the Male segment 1012. Exit and entry rules are then used to determine how to accommodate an appropriate update to the segment. The entry may be accommodated by taking the union of the existing membership in Male 1012 with the membership list in the entry membership list for Male. The exit may be accommodated by removing from the existing

membership in Male 1012 those audience members listed in the exit membership list (actual removal, of course, would only be applicable for those present prior to the recalculation).

[0098] For ease of discussion, focus is made on incremental profile data as it relates to Gender, but the principle of exit and entry can apply to any segment including but not limited to Visited News, HHI and others.

[0099] Incremental profile data based recalculation also propagates through the hierarchy.

This may be variously arranged, again depending upon exit and entry rules, which in turn depends upon the logical relationships of the segments. For a dependent (child) segment resulting from an AND operation such as Males who Visited News 1020, this may comprise repeating application of the above-described entry and exit membership lists for "Male" to the segment Males who Visited News 1020 in a similar fashion. That is, the entry membership list for Males would be added to the Males who Visited News 1020 segment, and the exit membership list removed. Alternatively, base segments Male 1012 and Visited News 1010 could be recalculated with their respective entry and exit membership lists, and then Males who Visited News 1030 could be calculated based upon the intersection of the updated versions of Male 1012 and Visited News 1010.

[00100] If desired, recalculation of a dependent segment could also be based upon a calculation based upon the updated parent segments. Specifically, the entry and exit 1032 functionality could be applied to the base segments, which could then be used to

[00101] FIG. 11A is a schematic diagram illustrating an example of processing 1100 data tables to manage and calculate segments according to another aspect of the present invention. The illustrated processing correlates with the segments that are defined in the example of FIGs. 10A-B. As described, the Extractor operates to collect information about numerous audience members and provides such information in the data warehouse. That information may be organized so that attributes corresponding to individual audience members may be identified.

The illustrated fact tables 1102a-d are a preferred technique for organizing the information as such. In one embodiment, each fact in a fact table is associated with an audience member using their unique identifier (PRID). A fact table contains all facts related to all users for a particular attribute. Accordingly, there is a section fact table that contains all section facts, an age fact table, a gender fact table, etc. Each row in a table represents a piece of data (characteristic or behavior) associated with only one audience member (more specifically one PRID).

[00102] As described attributes may involve characteristics such as age and gender as well as behaviors such as the number of times that the audience member has visited a particular section (News, Sports, etc.). At times, an attribute may be determined by looking at multiple pieces of information. Thus, while gender may be a simple determination of whether gender="male", an attribute that includes frequency information such as how many times an audience member visited a particular section may involve counting the number of entries in a fact table for the audience member. This counting may also be constrained to those entries falling within a particular time period.

[00103] Various alternatives may be used to provide the functionality of the fact tables, including different organization of the information. For example, the system may alternatively construct a table that provides a listing of attributes for a user identified by a unique PRID. This would result in a number of fact tables respectively corresponding to unique audience members identified by their PRIDs.

[00104] As previously described, the Segment Manager accesses the information stored in the data warehouse and maintains segment definitions, such as those input by the user seeking certain audience segments. A given segment is calculated by determining which audience members have the attribute for the given segment. According to this aspect of the present invention, the association of audience member identifiers to attributes and hierarchical logical

tree based segment architecture accommodate very efficient calculation (and recalculation) of segments.

[00105] A first level of processing 1104 may be used to calculate base segments. This is done by identifying the attribute for a base segment and then determining the audience members (or more particularly the listing of PRIDs) that have that attribute. Presume that segment 1.1 is the “Visited News” segment (see FIG. 10A). In this instance, the Segment Manager examines the fact tables and collect the PRIDs for those fact tables that contain this attribute. As indicated in segment table 1106a, this may result in a determination that PRIDs 1, 2, 4, 6, and 7 have the given attribute. The listing of PRIDs in a segment table may also be referred to as the “membership list” for the given attribute/segment. Again, there may be millions of members in a segment, the limited listings are used for ease of illustration.

[00106] The segments may also be identified by identifiers (SEGIDs) in lieu of the words and phrases that identify them. Thus associating identifiers SEGID_{x,x} with the noted PRIDs efficiently identifies the audience members with the attribute for computational purposes. Each segment may be organized in this fashion.

[00107] Continuing with the example, segment 1.2 may correlate to the attribute “Male”. Audience member PRID₁ is identified as male, and is listed in the segment table for segment 1.2, but PRID₂, identified as female, is not. The table 1106c for segment 1.3 (HHI > \$100K) includes both of those PRIDs. Again, segment tables for each of the segments may be provided, for x base level segments (1106a-d).

[00108] A next level of segments may then be calculated 1108 from the base segments. This aspect of the present invention accommodates efficient determination of further levels of segments through application of various Boolean operations to the existing segment tables. For example, Segment 2.1 may have been defined as “Visited News” AND “Male”. This is accommodated by determining the intersection of the PRIDs in those two segment tables

(1106a, 1106b). As illustrated, the segment table 1110a for segment 2.1 thus includes PRID₁, PRID₄, and PRID₆ since those identifiers appeared in both of the two base segment tables.

Table 1110a thus lists audience member identifiers for the males who have visited News. Once again, any number of segments may be calculated 1108 at this level, denoted as tables for segments 2.1 through 2.y (1110a-b).

[00109] Still further calculation 1110 accommodates determination of the next level of segments. Segment 3.1 (“Males who have visited News with HHI > \$100K”) correlates to a combination of Segment 2.1 (Males who have visited News) and Segment 1.3 (HHI > \$100K). Again, the logical AND implements the intersection of the relevant segment tables, which results in listing PRID₁ and PRID₄ as belonging to segment 3.1, per segment table 1114a. Any number of z segments may be calculated 1112 (segment tables 1114a-b).

[00110] The segment tables are the membership lists for their respective segments, and may be updated accordingly responsive to segment recalculation upon receipt of incremental profile data as previously described. FIG. 11B illustrates how the segment tables are updated responsive to recalculation based upon receipt of incremental data. Here, entry and exit is accommodated by tables containing membership lists, or entry tables and exit tables. As previously described incremental profile data (denoted respectively as fact tables 1102a’-d’) is received, and entry and exit tables are built based upon such data. FIG. 11B illustrates how the information in the entry and exit tables is useful for recalculating segments. Suppose that the entry table for the “Males” Segment 1.2 includes PRID₇ and the exit table for the same segment includes PRID₄. Application of the exit table would prompt PRID₄ to be removed from “Males” Segment 1.2 (as denoted by cross-hatching). Application of the entry table would cause PRID₇ to be added to the segment (as denoted “entry”). The membership of dependent segments is also updated according to the previously described logic. That is, because PRID₄ is no longer a member of Males Segment 1.2, it is also removed from dependent segment Males who have

Visited News 2.1. Continuing to the next level of dependency PRID₄ is removed from Segment 3.1, but PRID₇ is not added because Segment 3.1 is an AND combination of Segments 2.1 and 1.3, and PRID₇ is absent from Segment 1.3.

[00111] Note that different logical combinations will prompt different application of entry and exit upon recalculation. Segment 2.1 is a logical AND of Segments 1.1 and 1.2; if it were a logical OR combination of those segments, then PRID₄ would not be removed unless it was also removed from Segment 1.1.

[00112] Another aspect of the present invention provides profile synchronization. People may access various computers throughout the day and week, such as a home computer, office computer, mall kiosk, or the like.

[00113] As described above, PRIDs are unique identifiers that are used to identify and gather data regarding unique audience members. In that regard, when a new visitor (*e.g.*, a woman using her office computer) to a web site is encountered, they are associated with the next available PRID (*e.g.*, PRID_A). Cookies implemented in conjunction with the visitor's browser then include the particular PRID_A and are used to collect profile data for that visitor. Later on, the same person may use her home computer to visit the web site. Presuming that the home computer has not been used to access the site, there will not be recognition that she is the same person, and a new unique PRID (PRID_B) will be generated and associated with her behavior and characteristics from that computer. There will thus be two separate sets of profile data that actually correspond, unbeknownst to the Audience Targeting System, to the same person.

[00114] Further, the person may use another computer (*e.g.*, mall kiosk) that accesses the web site, and yet another unique PRID_C may be issued. This is problematic in two ways. First, it creates a third separate PRID for activity corresponding to the same person. Also, the mall kiosk (or even home and office computers) may be used by multiple people. Even though

multiple different people are using the computer and engaging in various behavior, it will all be tracked as PRID_C.

[00115] Still another problem is potential deletion of cookies. Continuing with this example, if this audience member deletes cookies on her office computer, then correlation with PRID_A is lost and she will be perceived as a new visitor on the next web site visit, prompting issuance of PRID_D in association with her office computer. This is problematic because the segments associated with PRID_D will not reflect information previously gathered in connection with PRID_A. Also, PRID_A will essentially become a defunct PRID, but will still be wastefully processed by the system.

[00116] FIG. 12 is a block diagram illustrating an example of an audience targeting system 1200 that includes profile synchronization 1260 according to another aspect of the present invention. Profile synchronization variously corrects and mitigates problems associated with these conditions. In one embodiment, the PRID is a system based identifier that uniquely identifies an audience member. An authoritative identifier (*e.g.*, a registration identifier) is also sought and maintained in association with a profiled audience member. An authoritative identifier may be identified in connection with some collected profile data. Maintenance of associations between authoritative identifiers and PRIDs allows such collected profile data to be properly associated with a particular audience member despite the absence of a PRID in the collected data. This functionality also accommodates the potential generation of multiple cookie based identifiers by a particular audience member. In contrast to the system identifier (PRID), which may also be referred to as an internal identifier, these cookie based identifiers are examples of external identifiers (XIDs). Maintenance of associations between each profiled audience member's PRID with one or more XIDs allows management of multiple external (*e.g.*, cookie based) identifiers in association with a particular audience member.

[00117] Before turning to a more detailed discussion of profile synchronization, it is noted that in embodiments of audience targeting that do not implement profile synchronization, the XID may essentially equate with the PRID for the purpose of audience member profile management. It is also noted that although cookie based XIDs are described, other external identifiers such as those that correlate to usage of a non-web device may also be implemented.

[00118] The Audience Targeting System 1200 includes a TE 1210, Extractor 1220, Segment Manager 1230 and Data Warehouse 1250. These elements are analogous to the commonly named elements in the previously described Audience Targeting System (800, FIG. 8) and need not be repeated with regard to the profile synchronization aspect.

[00119] As with the previously described system, the Audience Targeting System 1200 and its components are illustrated collectively, but may be provided individually and separately if desired. The functionality of the Profile Synchronization module 1260 is preferably provided by software that may be executed on any conventional processing system. In that regard, the audience targeting system 1200 (or any sub-module) may in turn be a component of a computer system containing a processor and memory. Although one modular breakdown is shown, it should be understood that the described functionality may be provided by greater, fewer and/or differently named components. Although a software embodiment is described, the functionality may also be provided as hardware or firmware, or any combination of software, hardware, and/or firmware.

[00120] The Profile Synchronization module 1260 includes an ID Management module 1262, an Authoritative ID Recognition module 1264, and an ID Storage module 1266 that in turn stores profile identifiers (PRIDs) 1268, REGIDs 1270, and XIDs 1272.

[00121] Profile synchronization entails a recognition that audience members, and the potential multiple identifiers that they may become associated with, may be associated with an authoritative identifier (ID). The Authoritative ID is in turn used to manage the multiple

identifiers as well as the profile data associated with the audience member. In one embodiment, the Authoritative ID is associated to registration (e.g., login credentials, REGID) for the user web site. For example, the web site may be The New York Times web site, which might require registration and login for usage of certain elements of the site.

[00122] The Profile Synchronization module 1260 implements PRIDs to uniquely identify audience members even as they generate multiple XIDs. In that regard, PRIDs may be regarded as system level, or more particularly Audience Targeting System 1200 level unique identifiers, and XIDs as audience member machine level based unique identifiers.

[00123] To accommodate the profile synchronization functionality, the ID Storage module 1266 stores the various ID information, including PRIDs 1268, REGIDs 1270, and XIDs 1272. The ID Management module 1262 organizes the issuance of and relationships between the various ID information. It accommodates this by associating the PRID for a particular user as uniquely identifying them on the system. This information may be stored along with other characteristics information such as the first date that the audience member was recognized by the system. Tables and the like may also be used to associate the audience member's PRID to the XIDs that are correlated to that audience member using profile synchronization, as well as to the REGID to accommodate recognition of audience members in conjunction with the Authoritative ID Recognition module 1264, which determines the presence of authoritative identification and communicates with the ID management module 1262 to ensure proper issuance of corresponding unique IDs.

[00124] The functionality of the Profile Synchronization module 1260 is further described with reference to the flow diagram of FIG. 13, which illustrates an example of a process 1300 for profile synchronization.

[00125] In support of the profile synchronization functionality, a new unique XID is associated 1302 with a first time visitor to the web site. If registration is applicable for the

session, then the REGID is associated as well. These functions are provided during regular browsing of pages and facilitated by the data agent as described above. Also in the described fashion, the data warehouse is populated with profile data corresponding to audience members. Unique REGIDs are thus also associated to respective sets of profile data along with the unique XIDs.

[00126] The profile data may be retrieved 1304 from the data warehouse in the previously described fashion. In embodiments using fact tables, this means that entries identifying both the XID and the REGID will be provided in association with the listed attributes. The fact table includes at least an XID, denoted particularly as XID_P in this example. A first determination 1306 is made as to whether a REGID is also included in the fact table. As described, the REGID is used as the authoritative ID. In its absence, the system seeks to process the data by attempting to associate the fact table with a PRID. As described, a list of XIDs is maintained in association with each PRID. This information is examined to see whether the particular XID (denoted XID_P) is found. If found, it is mapped to at least one PRID. It may be possible that an XID is mapped to multiple PRIDs. In that case the system may choose a random PRID, the first one found, or use any algorithm to select one. It should be noted that fact tables may be variously organized to provide this functionality. In one example of this the different attributes (Section, Age, Gender, Referrer, etc.) may each have a different table where a particular value is associated to a particular profile via the PRID.

[00127] With profile synchronization, the PRID uniquely identifies audience members for the purpose of segmenting. Accordingly, when it is determined 1308 that a particular PRID is associated with the particular XID_P , segments are calculated 1310 associating the attributes in the fact table to that particular PRID. If a PRID is not determined 1318 to be associated with XID_P , then a new $PRID_Q$ is issued 1312. In conjunction with that, XID_P is mapped to $PRID_Q$, and segments are calculated accordingly.

[00128] If it is determined 1306 that a REGID is present in the fact table, such is construed as the authoritative ID. This may be the first instance that the system sees a particular REGID, in which case a PRID is assigned (denoted $PRID_R$) and mapped to the REGID (1316).

[00129] If it is determined 1314 that there is already a PRID associated with the particular REGID (i.e., not the first instance of seeing REGID), then the particular PRID (the unique PRID number for that audience member) is associated to the fact table attributes and corresponding segments. Additionally, if such is not already the case, XID_P is included 1318 in the list of XID numbers that the system has associated to the particular PRID.

[00130] If desired, the segment manager may also segregate segments for an audience member using the XID list. For example, a particular audience member may have two XIDs associated to their unique PRID. One XID may correspond to his home computer and another XID may correspond to his work computer. Although the system will (through connection to the authoritative ID as described above) conclude that he is the same person and that all of the activities from both computers could be commonly segmented under the unique PRID, the listing of XIDs in association with that PRID allows the system to generate separate segments if desired. This may in fact be desirable to certain users of the Audience Targeting System since in some instances an audience member may have separate home and office personas in terms of computer usage and desired ad exposure.

[00131] FIGs. 14A-B are schematic diagrams illustrating an example of an Audience Matching Network 1420 according to an aspect of the present invention. The previously described audience targeting systems accommodate the definition of audience segments, the collection of profile data and corresponding determination of membership in those segments, and the delivery of content to audience members falling within such segments. A number of audience targeting systems (ATS) 1440, 1442, 1444 may respectively correspond to domains in which the process of defining segments and delivering content is carried out. Particularly, ATS

1440 may correspond to “*a.com*”, ATS 1442 to “*b.com*” and ATS 1444 to “*c.com*.” That is, visitors to *a.com* are audience members that may be targeted with advertising as they navigate among web pages in the *a.com* domain, via the functionality of the audience targeting system (here, ATS 1440) as previously described in detail.

[00132] According to one aspect, the present invention provides a network for matching an audience with deliverable content, which may be referred to as an Audience Matching Network (AMN) 1420. The AMN 1420 is an audience-centric network that allows advertisers to use behavioral targeting in combination with demographic data to reach defined audience segments of significant size. Preferably, the demographic data will be non-personal data. In some embodiments, personal data may be used, as constrained by audience member consent, contractual, and/or legal requirements.

[00133] The profile data for audience members may be collected across numerous sites having the audience targeting functionality (e.g., ATS 1440-1444). Each of these sites offers potential segments and population of such segments with members in their respective domains. These domains may also be referred to as “local” domains. These domains may, for example, each correspond to a separate publisher. As is well known, publishers may display online advertisements on pages, and advertisers typically pay publishers to place these advertisements.

[00134] The AMN 1420 has a domain (e.g., *audiencematchingnetwork.net*, or *amn.net*) that is common to the entire network of sites, which may be referred to as the network domain. The AMN 1420 presides over network level segments that comprise the various segments respectively populated by the ATS 1440-1444. As such, the AMN1420 is able to organize and manage segments based upon attributes that collectively traverse the network, and that otherwise would not have been identifiable by individual sites.

[00135] With continued reference to FIG. 14A, individual ATS may define segments and collect profile data for audience members in the relevant domain. For example, ATS 1440 may uniquely profile *a.com* audience member “123”, under a unique identifier denoted as PRID₁₂₃.

[00136] In the illustrated example, a “Traveler” segment may correlate to audience members who have been determined to be interested in travel, such as by their having visited a travel-related page in the domain. This particular segment in the *a.com* domain is denoted and managed as segment #100. The audience member managed under PRID₁₂₃ is determined to be a member of segment #100 (“Traveler”).

[00137] It should be noted that this example is illustrative only. As described above in connection with the ATS and segment management features of related inventions, segments may be variously defined according to characteristic and behavioral attributes, including but not limited to gender, section visited, HHI and others. It is also noted that data collection may not be limited to segment data, but may be any data points that a publisher allows the AMN 1420 to collect.

[00138] Still referring to FIG. 14A, in the *b.com* domain, ATS 1442 defines a “business” segment #200 and through collection of profile data within the domain, determines that the audience member uniquely identified in *b.com* as PRID₄₅₆ is a member of that segment. Further, in the *c.com* domain, ATS 1444 defines another segment #300 and determined that its uniquely identified audience member PRID₇₈₉ is a member of that segment.

[00139] Each of these audience members may be targeted for the delivery of content within respective domains. According to this aspect of the present invention, the AMN 1420 is able to recognize that an audience member that is (separately) uniquely profiled in different domains is in fact the same audience member, and to assemble segments that collectively traverse the numerous domains in the network, so as to determine that such an audience member is a member of a complex segment based upon information collected in the local domains.

[00140] The AMN 1420 uses a unique identifier referred to as a network PRID (NPRID) to manage the unique identification of audience members at the network level and to determine segment membership. The AMN 1420 also manages network level segments. For example, a “Business Traveler” segment may be managed as segment #5000 by the AMN 1420, and may be mapped to the segments (#100, #200) respectively defined in one or more of the domains in the network. The segment “Business Traveler” #5000 is shown for ease of discussion and to illustrate how the AMN 1420 may define “complex” segments that result from a combination of segments defined in separate local domains in the network.

[00141] The “Business Traveler” segment may be variously calculated, but preferably may comprise individual network level segments “Business” and “Travel” that are hierarchically organized. The determination of membership within base level segments and higher level segments that combine such base level segments may be accommodated through such organization of the segments. Also, membership tables can be used to calculate and recalculate segment membership using the NPRID identifiers at the network level.

[00142] As an alternative to having different segment identifiers for local and network domains, it may be more efficient to have a segment definition scheme that is universal to the network and local domains. Such a scheme would use “global” segment identifiers. Thus, for example, both the *a.com* ATS and the AMN 1420 may define segment #100 as “Travel”. This avoids mapping segments. As another alternative, the local domains may collect behavioral and characteristic information for passage to the network domain. The network domain would then collectively have the behavioral and characteristics information upon which targeted delivery of content may be based. The network designer is free to establish the system as desired.

[00143] The NPRID may be assigned the first time that an audience member is established with any one of the individual domains in the network, and cookie information may be used by the AMN 1420 to recognize a profiled audience member for future visits, even if such visits are

to sites in other domains in the network. Specifically, presume that a particular audience member visits *a.com*, and that the visit is the first of any site in the network. In connection with this, the particular audience member may be assigned PRID₁₂₃ by ATS 1440. In conjunction with this, the ATS 1440 (or, more particularly, the data agent associated with ATS 1440) directs the particular audience members browser to the AMN 1420, which assigns a unique NPRID (e.g., NPRID₁₀₁₁₁₂) to the particular audience member. In connection with this, the particular audience member browser may be provided with cookie information that identifies the particular audience member as PRID₁₂₃ in *a.com* and NPRID₁₀₁₁₁₂ in *amn.net*. Subsequently, the audience member may visit *b.com*, and be assigned PRID₄₅₆ by ATS 1442. However, when the browser is directed to AMN 1420, the AMN 1420 recognizes the audience member NPRID₁₀₁₁₁₂ based upon the cookie information and does not assign a new number. The AMN 1420 may, however, retain a listing of PRIDs corresponding to an NPRID. The cookie information may be variously organized. One example uses a local cookie corresponding to the local domain and a network cookie corresponding to the network domain.

[00144] The NPRID, in turn, is used to manage network segment membership and other network profile data values. Thus, for example, NPRID₁₀₁₁₁₂ would appear in the membership listing (e.g., tables) for the network level segments “Business” and “Travel”. A combination of these segments would provide a “Business Traveler” segment that would have NPRID₁₀₁₁₁₂ as a member. Also in connection with the collection of data and determination of segment membership, the AMN 1420 includes a targeting engine that, like the targeting engine in the ATS, may set segment cookies that identify the segments to which a particular audience member belongs based upon any collected network profile data. In this instance, such segment cookies are set in the *amn.net* domain.

[00145] FIG. 14A illustrates, among other things, the data collection role for members of the audience matching network. Another role that is provided in connection with the AMN

1420 is an “Ad Serving” role, illustrated in FIG. 14B. Although data collection may be performed by formal partners that are part of the audience matching network, the serving of ads to audience members who have visited sites in the audience matching network is not necessarily limited to those formal partners. “Non-Partner Site” *d.com* 1460 illustrates an example of an ad serving site, which serves audience matching network ads but does not collect data. Of course, partner sites may also fill the Ad Serving role.

[00146] In addition to allowing non-partner sites to serve AMN ads, the AMN 1420 accommodates the indication whether an audience member currently has any values in the network segment cookie that may be targeted against. The indication may be in the form of an AMN cookie (AMNC) with a Y/N value, where a “Y” indicates that there are values and an “N” indicates otherwise.

[00147] The serving of ads may be as follows. Someone visits 1480 the AMN Ad Serving Site 1460 (in the “*d.com*” domain). The AMN Ad Serving Site will likely incorporate an ad server (the ad server for *d.com*) to provide advertisements in connection with pages requested by the visitor, and as such the visitor’s browser is redirected 1482 to the *d.com* ad server (1484). An initial determination is made whether AMNC is set to “Y” in connection with the visitor’s browser. If this is not the case (or the AMNC Y/N is completely absent), then the *d.com* ad server serves non-AMN-network ads to the visitor’s browser in convention fashion. However, if the AMNC is set to “Y”, then the browser is redirected 1486 to the ad server in the *amn.net* domain (1488). There, the visitor is recognized as, say, NPRID₁₀₁₁₁₂ and is served 1490 ads appropriate for the segments in which NPRID₁₀₁₁₁₂ is a member.

[00148] FIG. 15 is a block diagram illustrating an embodiment of an AMN System 1500 that includes a targeting engine (TE) 1510, Extractor 1520, Segment Manager 1530, Audience Member Management 1540, Content Delivery and Management 1550, and Data Warehouse 1560 modules.

[00149] The AMN System 1500 and its components are illustrated collectively for ease of discussion. The various components and corresponding functionality may be provided individually and separately if desired, such as by different servers or agents that are assigned to the functionality of one or more of the components.

[00150] The functionality of the AMN system 1500 is preferably provided by software that may be executed on any conventional processing system, such as those previously named or others. In that regard, the AMN system 1500, or individual elements thereof, may in turn be a component of a computer system containing a processor and memory. Although one modular breakdown is shown, it should be understood that the described functionality may be provided by greater, fewer and/or differently named components. Although a software embodiment is described, the AMN system 1500 may also be provided as hardware or firmware, or any combination of software, hardware, and/or firmware.

[00151] As with the previously described analogous component in the ATS, the TE 1510 accommodates the assignment and coordination of unique identifiers corresponding to individual audience members in conjunction with the collection of data and the setting of cookies to support such collection of data, and, ultimately, the delivery of targeted content to the audience members.

[00152] With the AMN System 1500, when an audience member logs onto a page for the first time in the domain, the TE 1510 places a cookie on their browser, which contains a unique identifier. In contrast to the TE for the ATS, the domain is the audience matching network, which comprises the local domains of the various partner sites. Accordingly, the unique identifier for the network is set the first time any page for any partner site is visited.

[00153] This network level unique identifier is referred to as an NPRID, as described above. Although individual local sites may manage a profile through the previously described PRID, the AMN System 1500 uniquely identifies an audience member network-wide through

the NPRID. Whenever a particular audience member with a previously assigned NPRID returns to any partner site (e.g., *a.com*, *b.com*, or *c.com* in FIGs. 14A-B), the unique NPRID is sent back to the TE 1510. Using the NPRID as a basis to identify the audience member, the AMN System 1510 can set a network level segment cookie, which is subsequently used for the delivery of targeted content to the audience member's device. The TE 1510 may also create logs of this activity, useful for performance and revenue determinations.

[00154] The issuance of NPRIDs may be provided in conjunction with an Audience Member Management module 1540, which includes an network profile module (NPM) 1542 and a participation verification module 1544. The NPM 1542 manages the issuance of NPRIDs and retains lists of PRIDs corresponding to each unique NPRID. It may also communicate with the Segment Manager 1530 and thereby retain the lists of segments to which an audience member defined by an NPRID belongs.

[00155] The participation verification module 1544 accommodates the management of the AMNC value, including setting of the AMNC and changes to the AMNC setting. The AMNC accommodates an indication whether the corresponding audience member has values that may be targeted against.

[00156] The Extractor 1520 for the AMN System 1500 is similar to that for the ATS (FIG. 8, element 820), functioning in the network domain rather than the local domain.

[00157] Particularly, the AMN System 1500 Extractor 1520 similarly uses a data agent that is employed to assist in the gathering of information from website visitors, again provided in the form of code that is added to those pages in connection with which data collection is sought. That data agent may be thought of as having a network component and several local components, corresponding to the local domains. The primary difference is that the local data agent component sends data within the local domain, whereas the network data agent sends data within the network domain. As with the previously described data agent, the code may have

header and function call portions that respectively identify the functions and variables that it needs to operate and ensure that all variables have been collected. The data agent may also be configured to produce log lines suitable for receipt and processing by the TE. Examples of parameters include the version of the data agent, the page referrer, the page URL, time information, and the NPRID. Furthermore, the data agent may be used to carry out the setting of cookies related to the described AMNC value for determining participation in the AMN.

[00158] The Segment Manager 1530 is preferably as previously described in connection with the ATS Segment Manager (830, FIG. 8), and thus similarly includes segment organization, console management, segment generation with calculation & recalculation, and reporting modules, which need not be re-described. Notably, the Segment Manager 1530 accommodates the definition and management of segments corresponding to audience members based upon characteristic and behavioral information. The segments may be organized according to a hierarchical logical tree based architecture that allows scalable segment management and accommodates incremental recalculation of segments. Membership lists may be used to determine which audience members are part of which segments, with calculation and recalculation of segments upon exit and entry being processed in the same fashion. The significant distinction between the ATS segment manager and the AMN Segment Manager 1530 is that, in lieu of using the PRID as the basis for determining membership in segments, the Segment Manager 1530 uses the described NPRID. As described above, the Segment Manager 1530 may map ATS segments to network segments. Alternatively, ATS in the local domains may merely provide characteristic and behavioral information to the Segment Manager 1530, which collects the information and manages the definition and population of segments in any fashion that is desired.

[00159] The Content Delivery and Management module (CDMM) 1550 allows the determination of which advertisements (or other content) are to be delivered to which network-

level audience segments. In the context of this aspect of the present invention, this may merely be an identification of which advertisement corresponds to which segment. This functionality, and the corresponding information, may be exported to the AMN Ad Server for efficient serving of ads to visitors of sites in the network.

[00160] In addition to identifying the association of advertisements to segments, the CDMM 1550 may manage a bidding process whereby advertisers bid on AMN audience segments. The CDMM 1550 may also apply revenue and performance based management of audience segments and corresponding accounting. These aspects are described further in connection with FIGs. 17 et al. below.

[00161] Finally, the data warehouse 1560 is populated with and stores the various profile data as previously described, but does so at the network level.

[00162] FIGs. 16A-B are event diagrams illustrating an example of a computer implemented process 1600 for matching audience members to deliverable content according to the present invention. The described components of the AMN System are segregated to give a further understanding of the process 1600. For the *a.com* and *b.com* domains, the site and ad server are illustrated, as are the segment manager and targeting engine components of the ATS. For the *amn.net* domain, the AMN ad server, AMN targeting engine and AMN segment manager are shown.

[00163] The process initiates with a visitor (aka an audience member) requesting 1602 a page from Site A. In this instance, the audience member has been previously engaged and assigned PRID₁₂₃, and the requested page is related to travel, or is among those pages deemed to indicate an current interest in travel. Accordingly, the data agent for the domain causes a communication to be sent 1604 to Segment Manager A, indicating that the audience member with PRID₁₂₃ has requested the relevant page. This is just an example of various behavioral and characteristic information that could be provided to a segment manager. In response, Segment

Manager A processes 1606 received profile data, resulting in the inclusion of audience member PRID₁₂₃ in the travel segment, denoted as segment #100.

[00164] The segment information is passed 1608 to the Targeting Engine A. The visitor's browser is also prompted to call 1610 the Targeting Engine A, which receives the PRID, determines 1612 which segments correspond to that PRID, and then sets 1614 a local cookie to include a reference to the segment. Here, PRID₁₂₃ is associated with segment #100 in the *a.com* domain, so Targeting Engine A acts accordingly. Data agent code on a page visited by an audience member may initiate the call 1610. In this context, it would be any page subsequent to the page that generated the aforementioned data.

[00165] The AMN data agent also prompts the visitor's browser to provide 1616 information to the AMN Targeting Engine. For this example, the audience member is assumed to have previously been identified by the AMN System and as such has already been assigned NPRID₁₀₁₁₁₂. The inclusion of this audience member in *a.com* segment #100 is among the provided information, which is passed 1618 to the AMN Segment Manager, which then maps 1620 the *a.com* segment #100 to the AMN segment for Traveler, denoted as segment #4000. As described previously mapping is optional, and other forms of local collection and network organization of behavioral and characteristics information that do not require mapping or even local segment definitions may be used.

[00166] Finally, the AMN Segment Manager reports 1624 back to the AMN Targeting Engine, specifically that NPRID₁₀₁₁₁₂ is a member of network segment #4000. The AMN Targeting Engine then sets 1626 two cookies. The local (*a.com*) cookie is set to include identification that AMNC=Y, and the (AMN) network cookie is set to indicate that the audience member NPRID₁₀₁₁₁₂ is a member of network segment #4000 (and any other network segments to which the audience member belongs).

[00167] In connection with visiting site A, the audience member browser is also directed 1628 to the *a.com* domain Ad Server A to for advertisements (these may be those populating the currently visited page). In connection with this, Ad Server A is configured to inquire 1630 whether AMNC is set to indicate that an AMN network ad should be served, with the “Y” value indicating such to be the case. Presuming that AMNC does indicate this, the Ad Server A redirects 1632 the visitor’s browser to the AMN Ad Server. (If AMNC=N, the Ad Server A would simply serve a local advertisement).

[00168] The previously described setting of the network cookie includes the identification of the audience member as NPRID₁₀₁₁₁₂ as well as membership in segment #4000, and this information is thus sent 1634 to the AMN Ad Server in connection with obtaining the network ad. The AMN Ad Server processes 1636 this information, which results in sending 1638 the advertisement associated to the traveler segment to the particular audience member.

[00169] Continuing with the description of the process 1600 in connection with FIG. 16B, the audience member may similarly visit the *b.com* domain. The sequence of recognizing the audience member in connection with behavior that suggests “business” interest is particularly described. As was the case with the visit to the *a.com* domain, here the audience member requests 1640 a page identified as noting an interest in business. The audience member is (again, previously) identified as PRID₄₅₆ by the *b.com* ATS and as such the local cookie notes this information. The *b.com* domain data agent causes a communication to be sent 1644 to Segment Manager B, indicating that this audience member PRID₄₅₆ has requested the business page, and the Segment Manager B processes 1646 the information accordingly, resulting in the inclusion of audience member PRID₄₅₆ in the business segment, denoted as segment #200.

[00170] The Segment Manager B passes 1648 the information to the Targeting Engine B. Again, the visitor’s browser is also prompted to call 1650 the Targeting Engine B, which

receives the PRID, determines 1652 which segments correspond to that PRID, and then sets 1654 the local cookie to include a reference to the segment (#200).

[00171] As with the previously described visit to *a.com*, the AMN data agent prompts the visitor's browser to provide 1656 information to the AMN Targeting Engine. The audience member has previously been identified by the AMN System and assigned NPRID₁₀₁₁₁₂. The inclusion of this audience member in *b.com* segment #200 is among the provided information, which is passed 1658 to the AMN Segment Manager. If necessary, the AMN Segment Manager maps 1660 the *b.com* segment #200 to the AMN segment for Business.

[00172] The AMN Segment Manager may also determine membership in segments that result from a combination of local segments. Here, membership in "Business Traveler" is determined based upon membership in "Traveler" from data collected in the *a.com* domain, and membership in "Business" from data collected in the *b.com* domain. This "Business Traveler" segment may also be organized and managed numerically, such as network segment #5000 as noted. Also in connection with the segment, a determination 1662 is made whether the segment is targetable.

[00173] The AMN Segment Manager similarly reports 1664 to the AMN Targeting Engine that NPRID₁₀₁₁₁₂ is a member of network segment #5000. The AMN Targeting Engine then sets 1666 the local and network cookies, with the local (*b.com*) cookie set to include identification that AMNC=Y, and the (AMN) network cookie set to indicate that the audience member NPRID₁₀₁₁₁₂ is a member of network segment #5000.

[00174] The audience member browser is also directed 1668 to the *b.com* domain Ad Server B for advertisements. Ad Server B is configured to determine 1670 whether AMNC is set to indicate that an AMN network ad should be served, and if so redirects 1672 the visitor's browser to the AMN Ad Server. The identification of the audience member as NPRID₁₀₁₁₁₂ as well as membership in segment #5000, is sent 1674 to the AMN Ad Server in connection with

obtaining the network ad. The AMN Ad Server processes 1676 this information, which results in sending 1678 the advertisement associated to the business traveler segment to the particular audience member.

[00175] According to another aspect, the AMN System facilitates performance based content delivery, as well as accounting for revenue corresponding to the delivery of advertising.

FIG. 17 is a block diagram illustrating an example of an AMN system 1700 that includes advertising revenue and performance management.

[00176] As with the previously described AMN system (1500, FIG. 15), the AMN System 1700 includes a targeting engine (TE) 1710, Extractor 1720, Segment Manager 1730, Audience Member Management 1740, Content Delivery and Management 1750, and Data Warehouse 1770 modules.

[00177] Except as noted below, these modules 1710-1770 are similar to the analogous modules in the previously described AMN System. Also, the AMN System 1700 and its components are again illustrated collectively for ease of discussion. The various components and corresponding functionality may be provided individually and separately if desired, such as by different servers or agents that are assigned to the functionality of one or more of the components. As with the previously described AMN system, the functionality of this AMN system 1700 is preferably provided by software that may be executed on any conventional processing system, but may be variously provided and/or modularized as desired.

[00178] The most significant different between this AMN System 1700 and that previously described is the inclusion of placement 1752, bid process 1754, and revenue and performance modules 1756 in the CDMM 1750. The placement module 1752 provides the basic capability for accommodating the placement of content, typically in conjunction with an AMN ad server, as previously described.

[00179] The bid process module 1754 accommodates the bidding on network segments by advertisers. Advertisers bid on available segments, to help ensure that their advertisements target audience members corresponding to a desired segment, when those audience members subsequently visit the sites in the AMN network, or those sites that partner with the AMN to place advertisements. The AMN Ad Server will preferably be prompted to select an ad with the highest bid for a segment that has not already its performance goals, has not exceeded its run dates and is not excluded by a publisher exclusion rule. An audience member may be part of more than one segment, so this process occurs across multiple segment targets. The process may also provide a ranking of advertisements. Pages often include multiple locations for advertisements, and both the ranking of advertisements for a given segment and the membership in more than one segment may contribute to which advertisements are placed on a page. The algorithm for placing advertisements may be variously altered as desired. For example, a randomized entry of non-highest bid advertisements may be used to allow them to be placed, which may result in determining that such an advertisement is a leading performer in a given segment.

[00180] The revenue and performance module 1756 includes a placement performance module 1758, an information influence determination (IID) module 1760, a revenue sharing module 1762, and a dynamic segmentation module 1764.

[00181] Before describing the individual elements of the revenue and performance module 1756 in more detail, reference is made to FIG. 18, which is a flow chart providing an overview of delivering advertisements with revenue and performance management by the AMN System. As described in detail previously, information is collected 1802 corresponding to audience members who are visitors to web sites that comprise a network of information collection domains. Based upon the collected information, it is then determined 1804 that an audience member is a member of a network segment. This involves a determination that the audience

member profile data evidences attributes that are defined by the network segment. Armed with the knowledge that the audience member belongs to the network segment, the delivery of an advertisement is accommodated 1806. As described above, this is based membership of the audience member in network segment. It is also, preferably, based upon performance criteria that are described further in connection with the revenue and performance module 1756 below. In one embodiment, the performance criteria are configured to benefit the publisher corresponding to the placement of the advertisement. Finally, revenue may also be allocated 1808 corresponding to the delivery of the advertisement based upon participation as a data provider in the collection of information.

[00182] Still referring to FIG. 17, the placement performance module 1758 generally determines which advertisements to deliver, based upon various criteria but preferably making publisher (web site) revenue paramount in placement determinations. The revenue sharing module 1764 allocates the sharing of revenue related to ad placement and related activities according to the role that a party provides. The information influence determination module 1760 works in conjunction with the revenue sharing module 1762 and helps in the allocation of revenue by establishing the influence that various different pieces of information is deemed to have, particularly with regard to how recent such information was collected. The dynamic segmentation module 1764 accommodates the creation of segments to maximize the revenue and yield from advertisements placed in connection with such segments.

[00183] In addition to the delivering advertisements to audience members belonging to network-defined segments based upon behavior and the desired target audience of an advertiser as described above, the placement performance module 1758 optimizes the delivery of advertisements from an existing 3rd party network. For example, a publisher may seek to optimize the delivery of all their inventory currently served by a 3rd party ad aggregator or optimization service (such as Google, Overture, Kanoodle, Advertising.com, etc).

[00184] In this regard, the AMN network acts as a gateway for publisher partners, proxying requests for ads on the publisher's pages, with the placement performance module 1758 choosing the correct ad for a given page/position based upon a set of criteria including maximizing revenue and fulfilling inventory commitments.

[00185] Two examples of maximizing revenue to a publisher include 1) choosing, from among the pool of possible ads, those ads for which the publisher is paid the most money (impression model) and 2) increasing the likelihood of generating a click-through for any given ad by providing additional behavioral targeting criteria not currently available to the 3rd party networks.

[00186] In the first example, the publisher provides the AMN System with its impression goals and the CPM (cost per thousand) for each of the potential ad deliveries/networks. The placement performance module 1758 then uses this information to direct the AMN Ad Server to choose, based upon current delivery metrics (*e.g.*, number of impressions previously delivered against a given campaign, CPM), the advertisement that will provide the most revenue to the publisher. This arrangement improves on the approach of simply serving the ad with the highest CPM in every instance, as each campaign has a maximum impression goal beyond which the advertiser or network will not pay for impressions.

[00187] In the second example, the AMN Ad Server has available the segment affiliation of the visitor and can choose, based upon the target criteria of the campaigns, the ads most likely to generate a click from the visitor, thereby maximizing the revenue for the so called pay-per-click campaigns. Click likelihood in this case may be determined in a number of ways. One example is to evaluate the click history of the audience member and see what ads s/he has clicked on in the past. Another example would be to look at the ads that have the highest click rate with respect to the target segment or other segments. There may also be a feedback loop, of clicks /

per impression / price per click to determine the value of an ad, in lieu of purely determining the ads most likely to be clicked.

[00188] These are just two examples of optimization that may be performed by the AMN System through the placement performance module 1758. Publishers and advertisers may provide various types of information for application to ad delivery opportunities as desired.

[00189] The IID module 1760 determines the influence that information has on an audience member's behavior. According to this aspect of the present invention, not all data is presumed to have equal value in influencing the behavior of an audience member. For example, some data may be more valuable in influencing behavior of one type (*e.g.*, auto purchase) while other data is more valuable for other types of behavior (*e.g.*, traveler). The IID module 1760 evaluates the relative impact that different pieces of data have upon the effectiveness of a particular segment. These effectiveness metrics can then be used for determining the priority of an ad to be served to an audience member (*i.e.* to maximize the revenue by serving the ads most likely to be clicked on by the audience member) and for the determination of distribution of revenue to data providers as described with regard to the Revenue Sharing module 1762, below.

[00190] Various factors that may go into determining the impact that individual data points have on the overall effectiveness of a segment, and may be configured as desired by the system designer. Examples of factors include recency, frequency, correlation analysis across segments with similar data values, data half-life analysis (described below), primacy (first site to provide data value), and intensity (which is similar to frequency but is measured as frequency per unit time).

[00191] The data half-life analysis refers to a process whereby an analysis is done of the overall effectiveness of data values over time to determine at what point a data value's effectiveness has been reduced by half. For example, it is well known that car buying behavior is relevant for less than 90 days as a consumer researches and then either purchases or abandons

the purchase process. In one embodiment, this is determined statistically by looking at the click-through rates on ads targeted at segments including particular data values and correlating this information with the age of the data provided. This is used to generate a curve where the effectiveness declines over time. This half-life analysis is then used to influence the value of a particular data value when evaluating a segment that contains that data value.

[00192] In addition, data provider metrics may also influence the value that a particular data point has in the overall effectiveness evaluation. For example, cars.com or Edmunds might be considered a more reliable indicator of auto buyer interest than similar content at USAToday.com or the Dallas Morning News website.

[00193] Finally, a further improvement on equal weighting of all data values is considering the value of the audience that is a part of a given segment in determining the effectiveness of the data in a segment. This is provided by analyzing behavior by a segment. If the audience included has a high proportion of audience members that have shown a willingness to click on ads, this information can be used to further enhance the effectiveness profile of the segment. Other metrics in this category include analysis of the average number of segments that the audience members in a given segment are a part of. This could show a more valuable audience because the diversity of data about this audience, leading to more "touch points" and hence, opportunities to present relevant advertisements.

[00194] The Revenue Sharing module 1762 accommodates allocation of revenue among relevant participants in the delivery of advertisements. Generally, the Revenue Sharing module 1762 builds upon standard practices in the industry where by two or more partners agree to split the revenue generated by an ad placement.

[00195] In one case, this involves inclusion of the publisher that provides the data about a targeted audience member and may thereby be referenced as a "data provider" in the revenue sharing arrangement. The data provider is compensated based upon the fraction of the data

used in the targeted delivery provided by the data provider. For example, if an audience member is classified into the network segment “Business Travelers”, and the data is provided in line with the example described above, then *a.com* and *b.com* would split the data provider revenue when a Business Traveler clicks on an ad, since they each contributed 1/2 the data involved in the segment that resulted in a click. In addition, the audience provider (publisher partner that served the ad that was clicked on) will also receive a percentage of the revenue.

[00196] The Revenue Sharing module 1762 also applies data value and audience value metrics, as described regarding the IID module 1760, to the revenue sharing arrangement. This accommodates rewarding the data providers that contribute the audience member or data value that had more impact on the successful conversion of the audience member (that is to say, the click event) with a higher percentage of the data provider portion of the revenue.

[00197] Finally, the Revenue Sharing module 1762 may also determine data provider value metrics in line with the information provided by the IID module 1760. In this scenario, the brand or nature of a data provider’s site will result in a higher or lower apportionment of the data provider revenue share based upon the value of the brands market awareness or authority of the data providers content.

[00198] The dynamic segmentation module 1764 generates segments based upon a correlation analysis. This may informally be referred to as a “people like this” approach. This allows an advertiser to identify a desired action or series of actions and tag the pages (on their site or others) that represent the desired action (purchase, visit of a promotional site, etc). Based upon the historical behavior of the visitors that take the desired action, the dynamic segmentation module 1764 correlates the profiles of visitors that have the desired behavior with all other profiles to find those visitors that have not taken the desired action but show a close affinity to those profiles that have taken the desired action. Consequently, the AMN Ad Server

may be instructed to target a promotional advertisement to those visitors that have not yet taken the action, with a high degree of likelihood that these users will also take the desired action.

[00199] FIG. 19 is a block diagram illustrating an example of an AMN system 1900 that includes universal profile synchronization (UPS) according to another aspect of the present invention. The AMN System 1900 includes a targeting engine 1910, Extractor 1919, Segment Manager 1930, CDMM 1950, and data warehouse 1960. Each of these modules are described in connection with the AMN system (1500) of FIG. 15 and need not be re-described in connection with this aspect of the present invention. In contrast to that system, the AMN System 1900 audience member management module 1940 is updated to include a Universal Profile Synchronization (UPS) module 1946.

[00200] Cookie blocking technologies have become an increasing problem for online publishers. While there are certainly many legitimate uses for this technology, the overly broad approach that most implementations take affects not only the troublesome adware and spyware that most users wish to block, but also the benign state management cookies that most online publishers rely on to manage their business and their relationship to the consumer.

[00201] As described above in connection with profile synchronization in the context of an ATS, publishers typically retain information that can be used as an authoritative identifier for a given audience member, for example registration login credentials. The publisher also has an interest in protecting the privacy of audience members, and to enforce the obligations in its privacy policy. The AMN System, configured to include the UPS module 1946, implements authoritative identifiers along with the collective information of the network to provide a robust network profile that can help overcome the problems associated with cookie blocking technologies. That is, an audience member can be reconnected with their data after cookies may have been deleted – or even if the audience member moves to a different client machine.

[00202] This dramatically simplifies the consumer profile management process for both publisher and consumer by creating a single authoritative source for site-specific data, as well as a clearing house for global, non-personally identifiable consumer data that may be shared across some or all publishers within the network.

[00203] The UPS functionality implements two primary requirements – the collection of specific data points associated with an anonymous profile at the publisher site and a network mechanism for identifying that visitor authoritatively upon visiting any site within the network.

[00204] For each publisher in the network, this entails code within the context of the publishers domain(s) that captures the identified data points when a visitor is consuming content on that site. The AMN system data agent, as previously described, captures the local segment affiliation from the segment cookie in the content domain and sends that information back to the AMN Targeting Engine. With the UPS functionality, the data agent is further augmented to know what data points each publisher has authorized for capture and where to find that data. The publishers content delivery engine -- via HTML tags within the content pages -- may explicitly pass this data to the data agent. The data agent configured as such is referred to as the UPS data agent (UPSDA).

[00205] Once the UPSDA has captured the data, it sends the data back to the UPS Data Collection Engine (UPSDCE). The UPSDCE is similar to the described Targeting Engine (for both the ATS and AMN Systems). However, the UPSDCE immediately captures the data points and stores them in a local database associated with the audience member's profile.

[00206] Preferably, the UPSDCE resides within a domain owned and managed by the provider of the audience matching network, in order to allow for the association of data to a global profile (NPRID), rather than just the local profile (PRID) of the individual publisher. The data also clearly identifies which publisher the data collected belongs to, in order to ensure that audience members' privacy is maintained by not inadvertently co-mingling different publishers

data. The assets of each publisher are also protected by not sharing this data with other publishers, unless specific agreements exist to share data within the network.

[00207] To properly maintain data integrity, it is useful to authoritatively identify a visitor within the network. There are a number of different ways to accomplish this. One approach is to utilize those publishers that have required registration and login to re-identify the visitor. Another approach is to use client side browser plug-in technology to store an authoritative identifier for that visitor. Either or both of these techniques may be used to assist in authoritative identification of a visitor.

[00208] Once a visitor is identified, the login credentials can be used to look up the visitor in the network repository (that is, the data set collected previously by the UPSDCE) and compare the current set of cookie values with those in the repository. If no cookies are currently set, or if the data values do not match, the visitor's cookies are updated to reflect the profile cookie data.

[00209] A visitor to the network need only be authoritatively identified once in order to authoritatively identify this visitor to all sites within the network, since the network identification can serve as a key into the profile data for that visitor across all sites.

[00210] A further service is the ability to share limited anonymous profile data across the network, based upon the publisher's willingness to opt-in certain information and also in compliance with the publisher and network privacy policies.

[00211] Finally, it is important that this data get updated on each page that a visitor views in order to ensure current and correct values, many of which are likely modified during a visitors session with the site.

[00212] FIG. 20 is a flow diagram illustrating an example of a process for UPS as implemented by the appropriately configured AMN System. The process 2000 generally arises in connection with the maintenance and profiles of audience members targeted for the delivery

of content. As described, this involves receiving 2002 profile data for a plurality of audience members corresponding to a network of local domains which may be referred to as audience member information collection domains. A given audience member has an associated set of profile data based upon information collected in the various local domains. The NPRID is used to maintain 2004 an association between the set of profile data and the profiled audience member. According to this aspect of the present invention, the NPRID is related 2006 to a set of information that facilitates the collection of profile data for the profiled audience member in the network of information collection domains. This information is preferably cookie related information that is used in connection with the collection of data in the network and various local domains, referred to previously as network cookie and local cookies. The NPRID is also associated with the previously identified authoritative identification for the profiled audience member.

[00213] In connection with a visit to any site within the network, an authoritative identification is received 2008. This information may be received in the absence of the NPRID. The authoritative identification identifies the profiled audience member in connection with activity, and is used 2010 to correlate the profiled audience member to the NPRID. In turn, the NPRID is associated to the cookie related information as described. This allows a comparison 2012 of the cookie information connected with the current activity with that stored in association with the NPRID. Such information can be used to update 2014 the cookie information in association with the audience member's browser, even if the cookies have been deleted between past profiling and the current browsing activity, or even if the audience member uses a different machine (if desired). Such updating may of course entail restoring the cookie information previously established for this particular audience member.

[00214] Thus embodiments of the present invention provide an audience matching network and related aspects. Although the present invention has been described in considerable detail

with reference to certain embodiments thereof, the invention may be variously embodied without departing from the spirit or scope of the invention. Therefore, the following claims should not be limited to the description of the embodiments contained herein in any way.

CLAIMS

1. A computer implemented method for performance based delivery of content to an audience member, the method comprising:

managing audience segments corresponding to a network of audience member information collection domains;

determining that a profiled audience member is a member of a network segment that includes audience members having a set of attributes corresponding to profile data collected in the network of audience member information collection domains; and

accommodating the delivery of an advertisement to the profiled audience member based upon membership in the network segment and performance criteria configured to benefit a publisher of content in which the advertisement is placed.
2. The method of claim 1, wherein the performance criteria includes maximizing revenue to the publisher.
3. The method of claim 2, wherein accommodating the delivery of an advertisement includes selecting an advertisement with the highest available payment to the publisher.
4. The method of claim 1, further comprising:

receiving from the publisher a set of delivery criteria corresponding to potentially deliverable advertisements; and

selecting an advertisement with the highest available payment to the publisher based upon a comparison of current delivery metrics to the set of delivery criteria.
5. The method of claim 1, further comprising:

optimizing the delivery of advertisement inventory served by a third party.
6. The method of claim 5, wherein optimizing the delivery of advertisement inventory

served by the third party includes providing information about membership in the network segment in connection with an advertisement serving function of the third party.

7. The method of claim 1, wherein the network segment is defined by examining the data profiles for audience members that have engaged in a desired behavior.

8. An apparatus for performance based delivery of content to an audience member, the apparatus comprising:

means for managing audience segments corresponding to a network of audience member information collection domains;

means for determining that a profiled audience member is a member of a network segment that includes audience members having a set of attributes corresponding to profile data collected in the network of audience member information collection domains; and

means for accommodating the delivery of an advertisement to the profiled audience member based upon membership in the network segment and performance criteria configured to benefit a publisher of content in which the advertisement is placed.

9. The apparatus of claim 8, wherein the performance criteria includes maximizing revenue to the publisher.

10. The apparatus of claim 9, wherein accommodating the delivery of an advertisement includes selecting an advertisement with the highest available payment to the publisher.

11. The apparatus of claim 8, further comprising:

means for receiving from the publisher a set of delivery criteria corresponding to potentially deliverable advertisements; and

means for selecting an advertisement with the highest available payment to the publisher

based upon a comparison of current delivery metrics to the set of delivery criteria.

12. The apparatus of claim 8, further comprising:

optimizing the delivery of advertisement inventory served by a third party.

13. The apparatus of claim 12, wherein optimizing the delivery of advertisement inventory served by the third party includes providing information about membership in the network segment in connection with an advertisement serving function of the third party.

14. The apparatus of claim 8, wherein the network segment is defined by examining the data profiles for audience members that have engaged in a desired behavior.

15. A computer implemented method for allocating revenue corresponding to the delivery of content to an audience member, the method comprising:

managing audience segments corresponding to a network of audience member information collection domains;

determining that a profiled audience member is a member of a network segment that includes audience members having a set of attributes corresponding to profile data collected in the network of audience member information collection domains;

accommodating the delivery of an advertisement to the profiled audience member based upon membership in the network segment; and

allocating revenue corresponding to the delivery of the advertisement based upon participation as a data provider in the collection of information.

16. The method of claim 15, wherein the network of audience member information collection domains includes a first domain and a second domain that is separate from the first domain, and the set of attributes includes a first attribute based upon information collected in

association with activity by the profiled audience member in the first domain and a second attribute based upon information collected in association with activity by the profiled audience member in the second domain.

17. The method of claim 16, wherein allocating revenue includes accounting for participation of a first data provider for providing the first attribute and of a second data provider for providing the second attribute.

18. The method of claim 15, further comprising:
determining an influence that information is deemed to have on audience member
behavior; and
accounting for the influence in allocating revenue corresponding to the delivery of the
advertisement.

19. The method of claim 18, wherein determining the influence includes an examination of a set of factors associated with the collection of the information.

20. The method of claim 19, wherein the set of factors includes at least one factor selected from the group comprising recency, frequency, and primacy.

21. The method of claim 19, wherein the set of factors includes a half-life analysis.

22. The method of claim 19, wherein the set of factors includes at least one factor selected from the group comprising recency, frequency, segment correlation, half-life analysis, primacy, and intensity.

FIG. 1




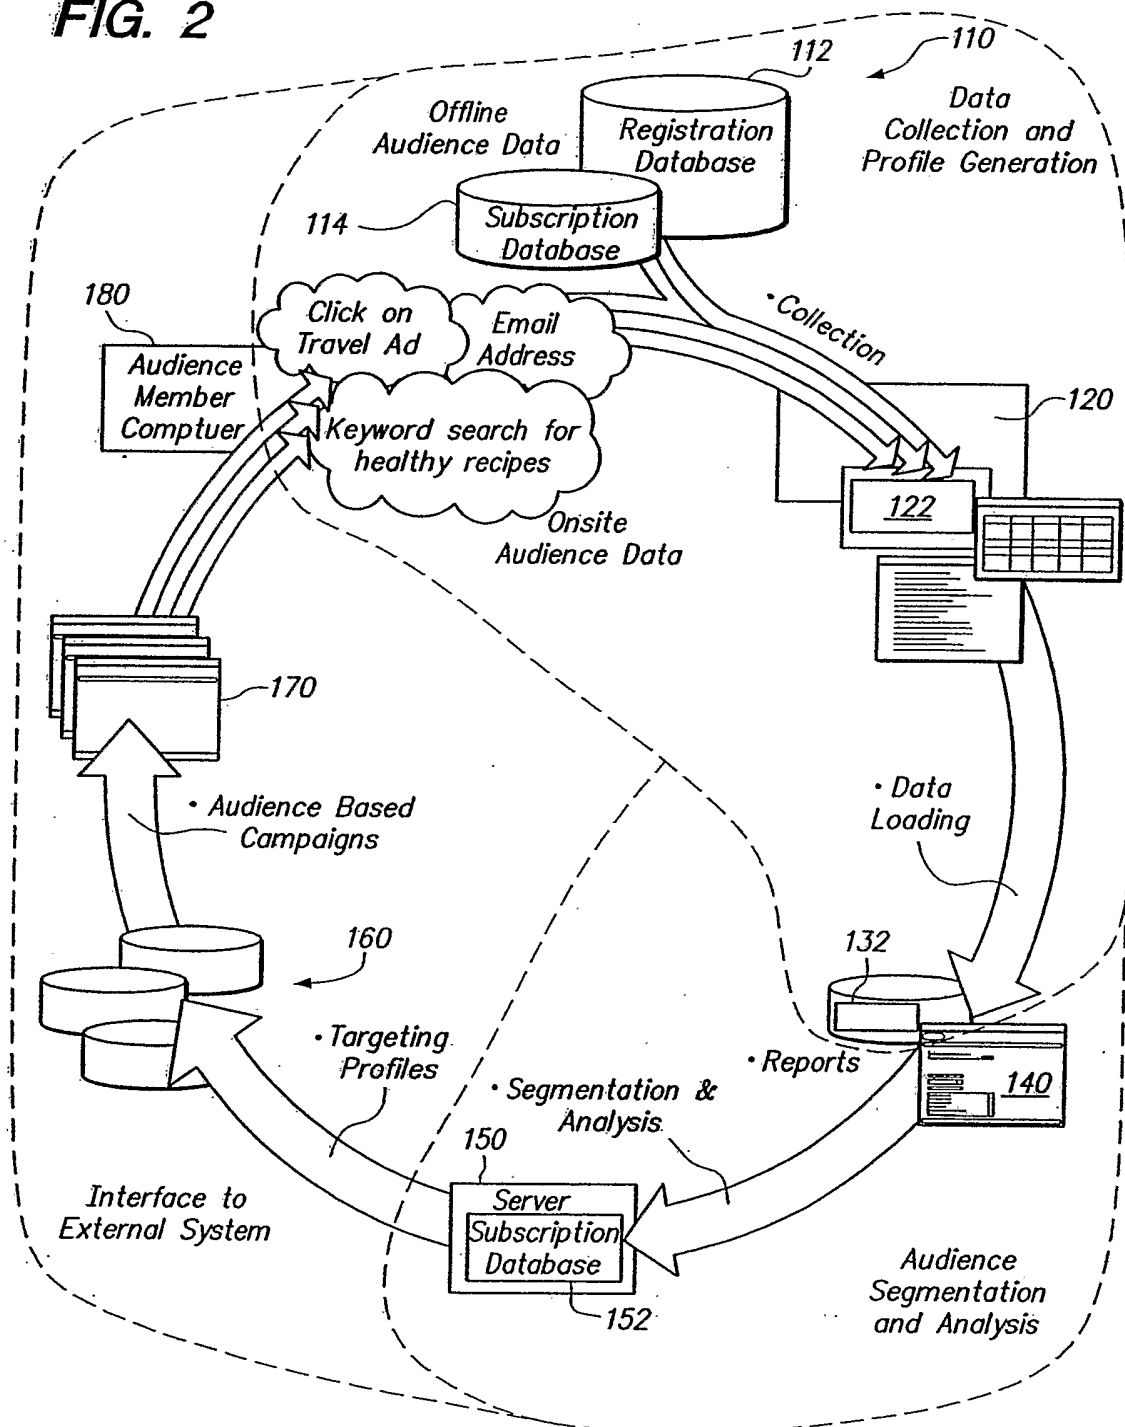
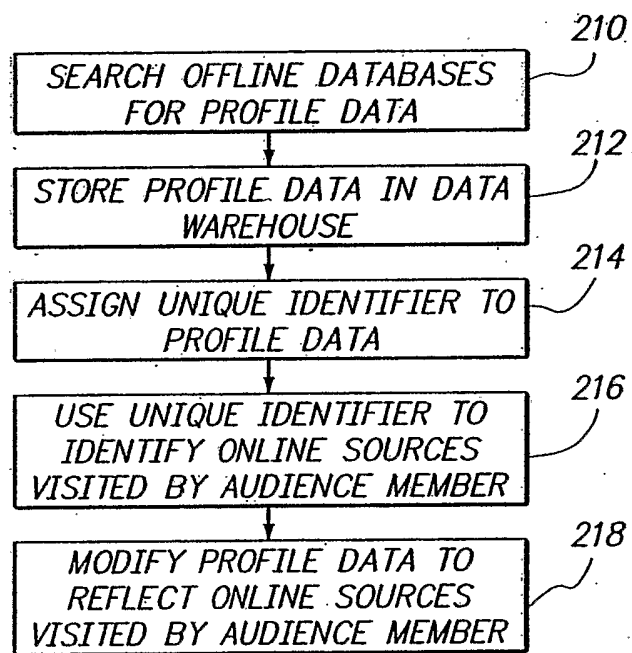
FIG. 2

FIG. 3





FIG. 5

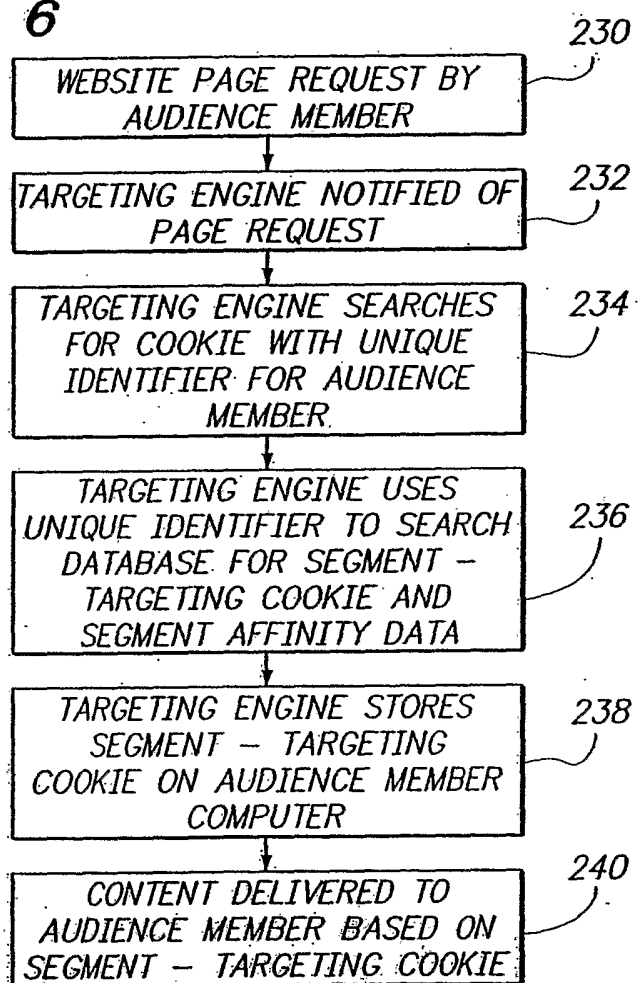
FIG. 6

FIG. 7





FIG. 8

FIG. 9A




FIG. 9B

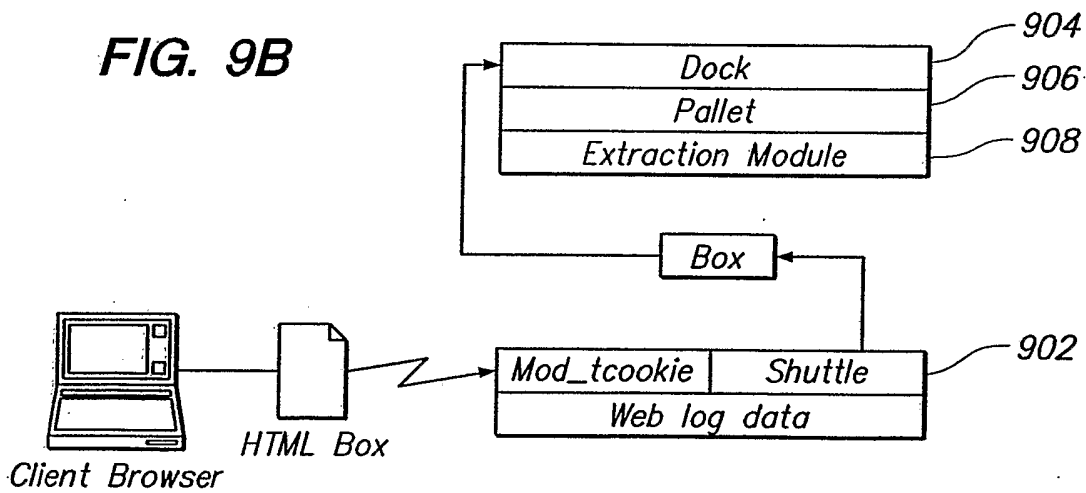


FIG. 10A




FIG. 10B




FIG. 11A

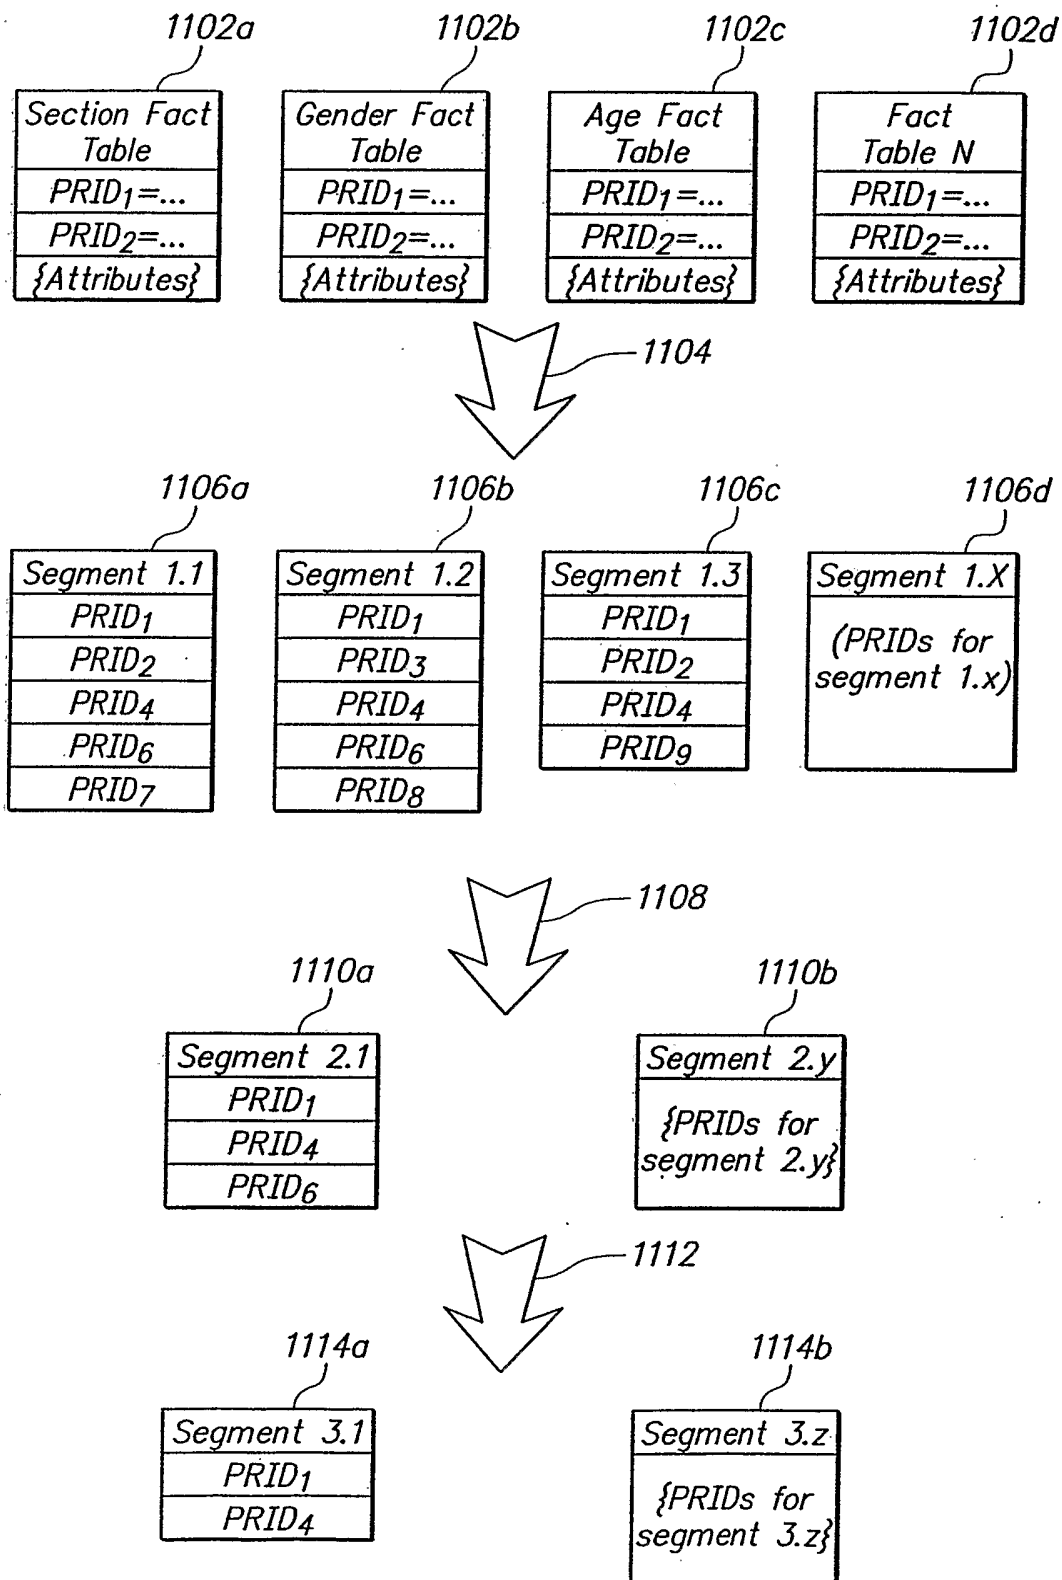


FIG. 11B





FIG. 12

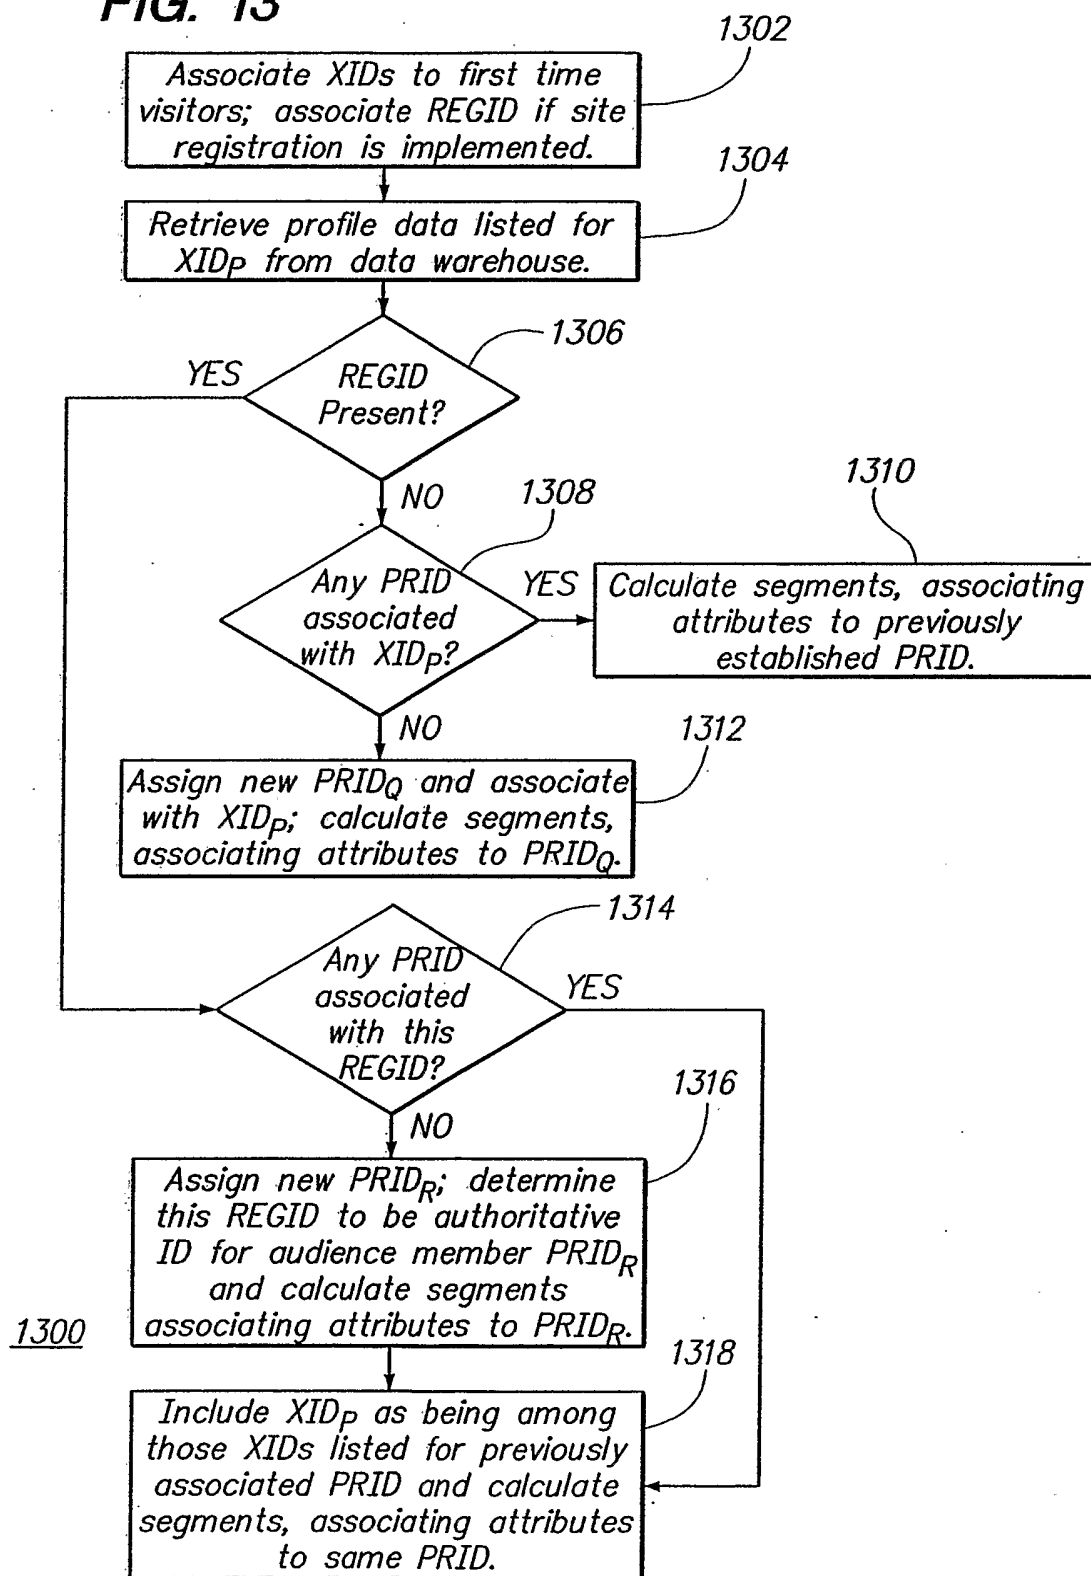
FIG. 13

FIG. 14A




FIG. 14B

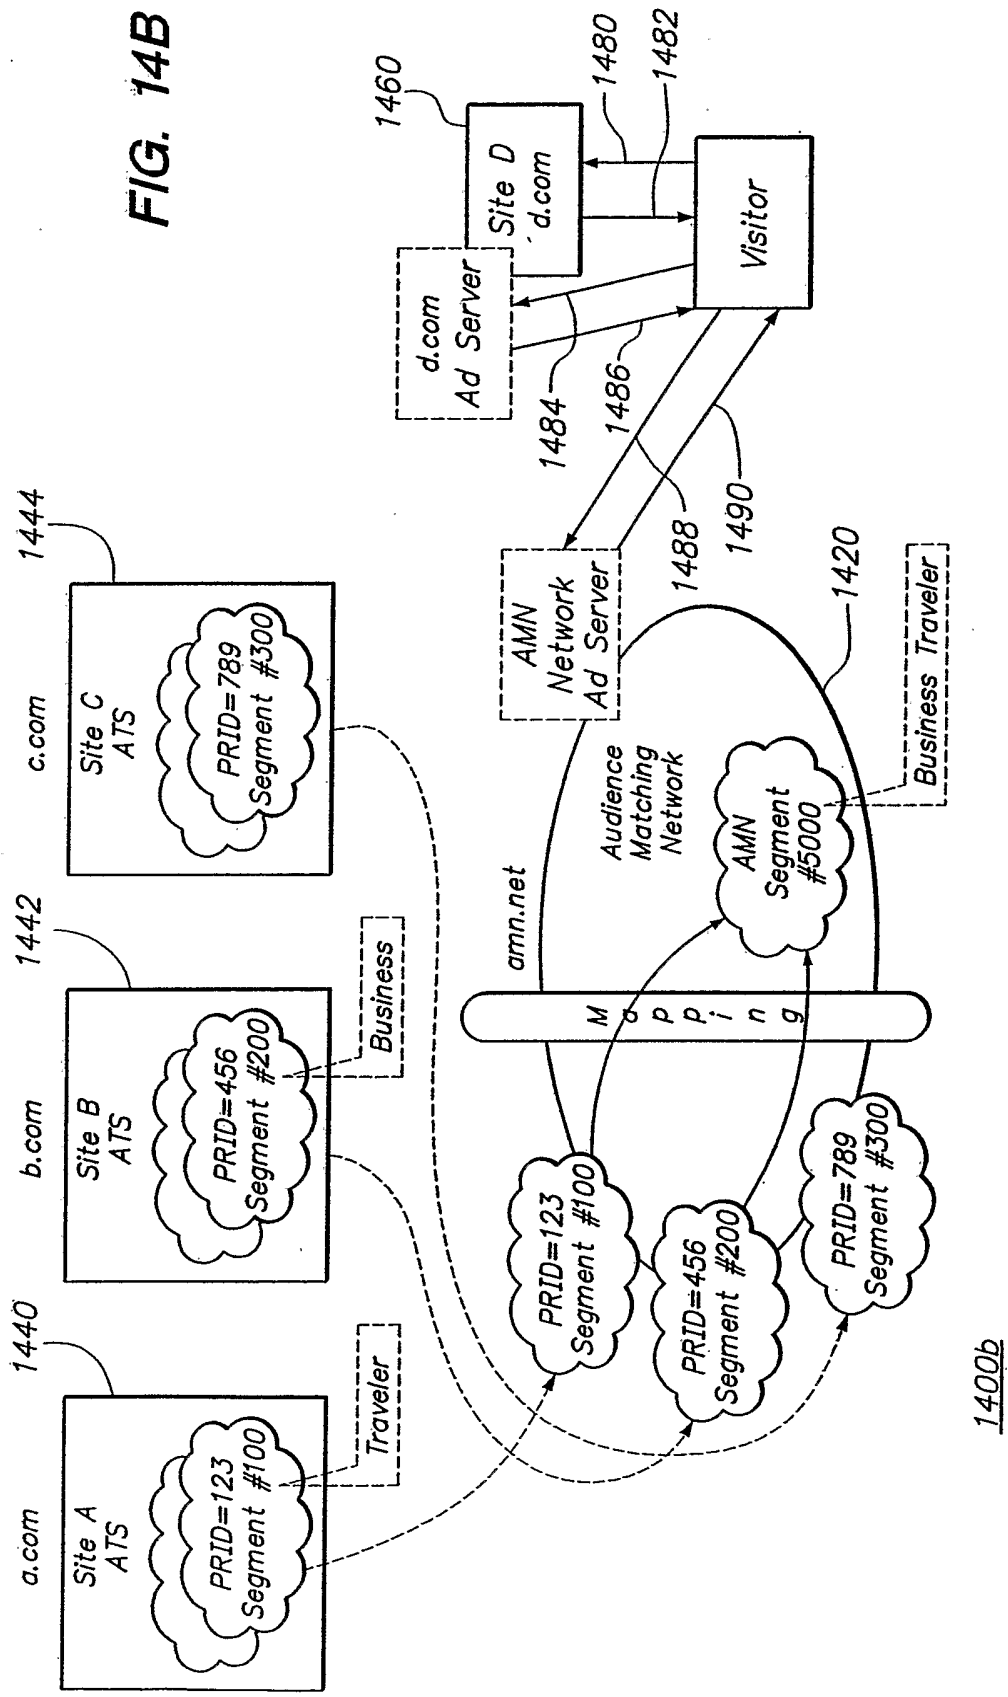



FIG. 15

FIG. 16A-1

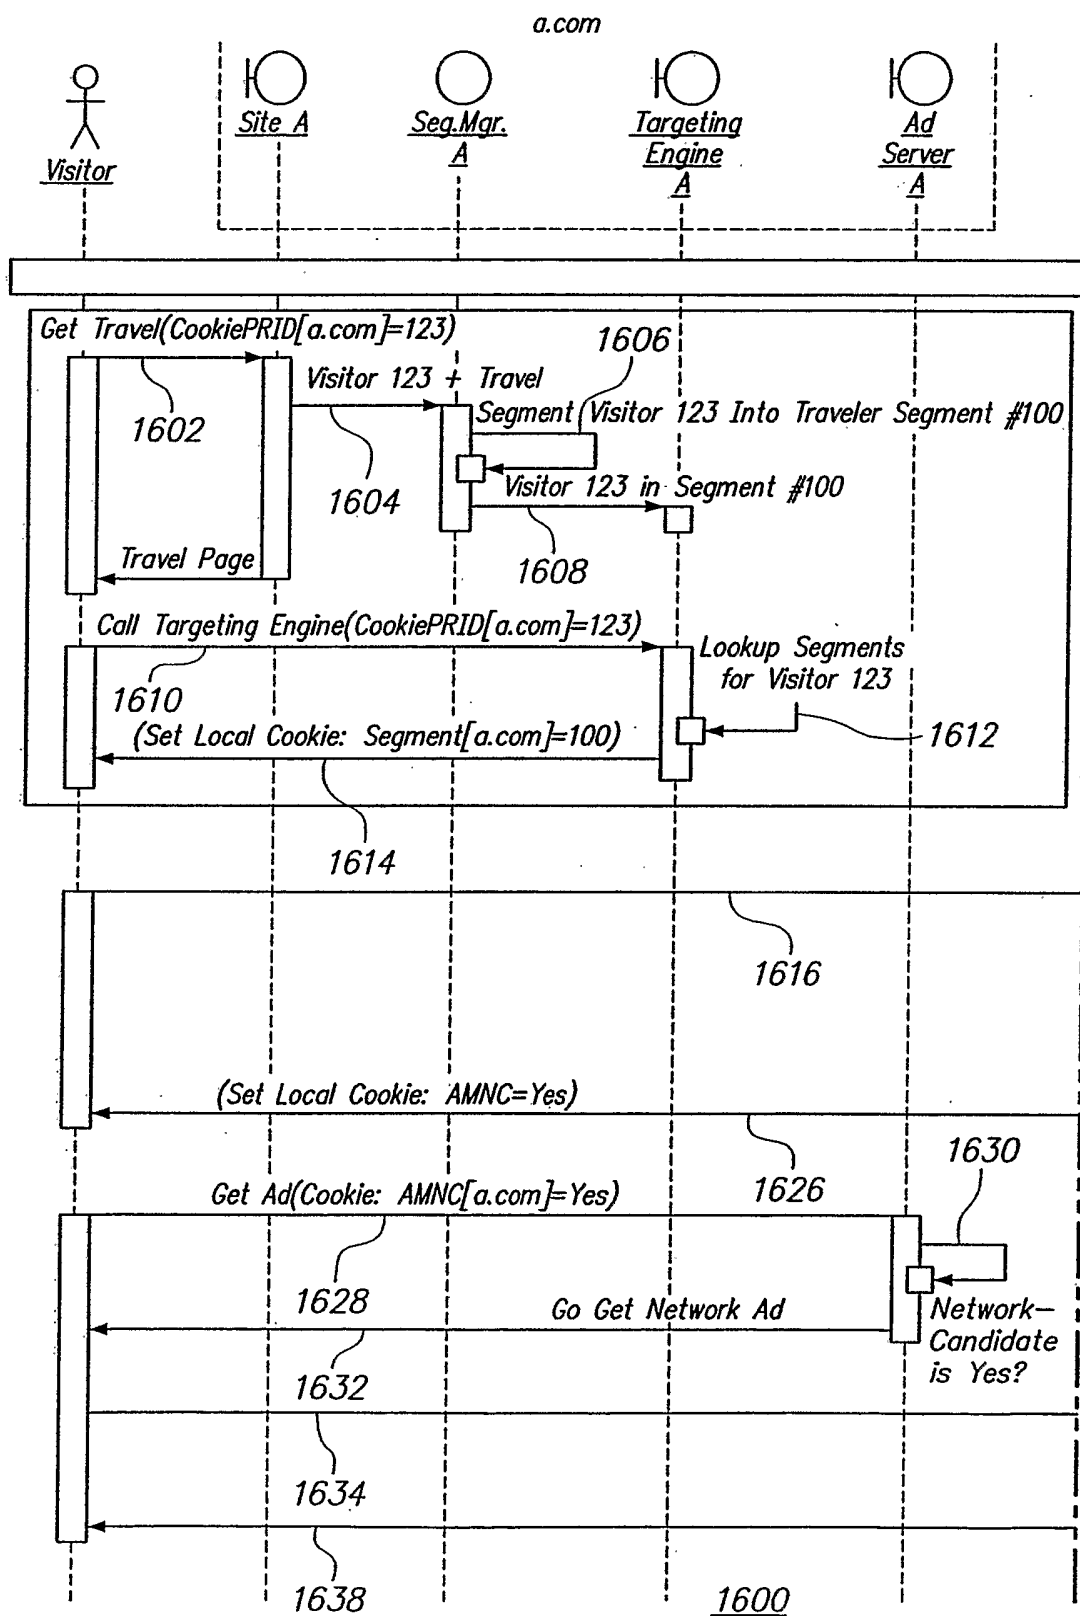


FIG. 16A-2




FIG. 16A-3




FIG. 16A

FIG. 16A-1	FIG. 16A-2	FIG. 16A-3

FIG. 16B-1




FIG. 16B-2

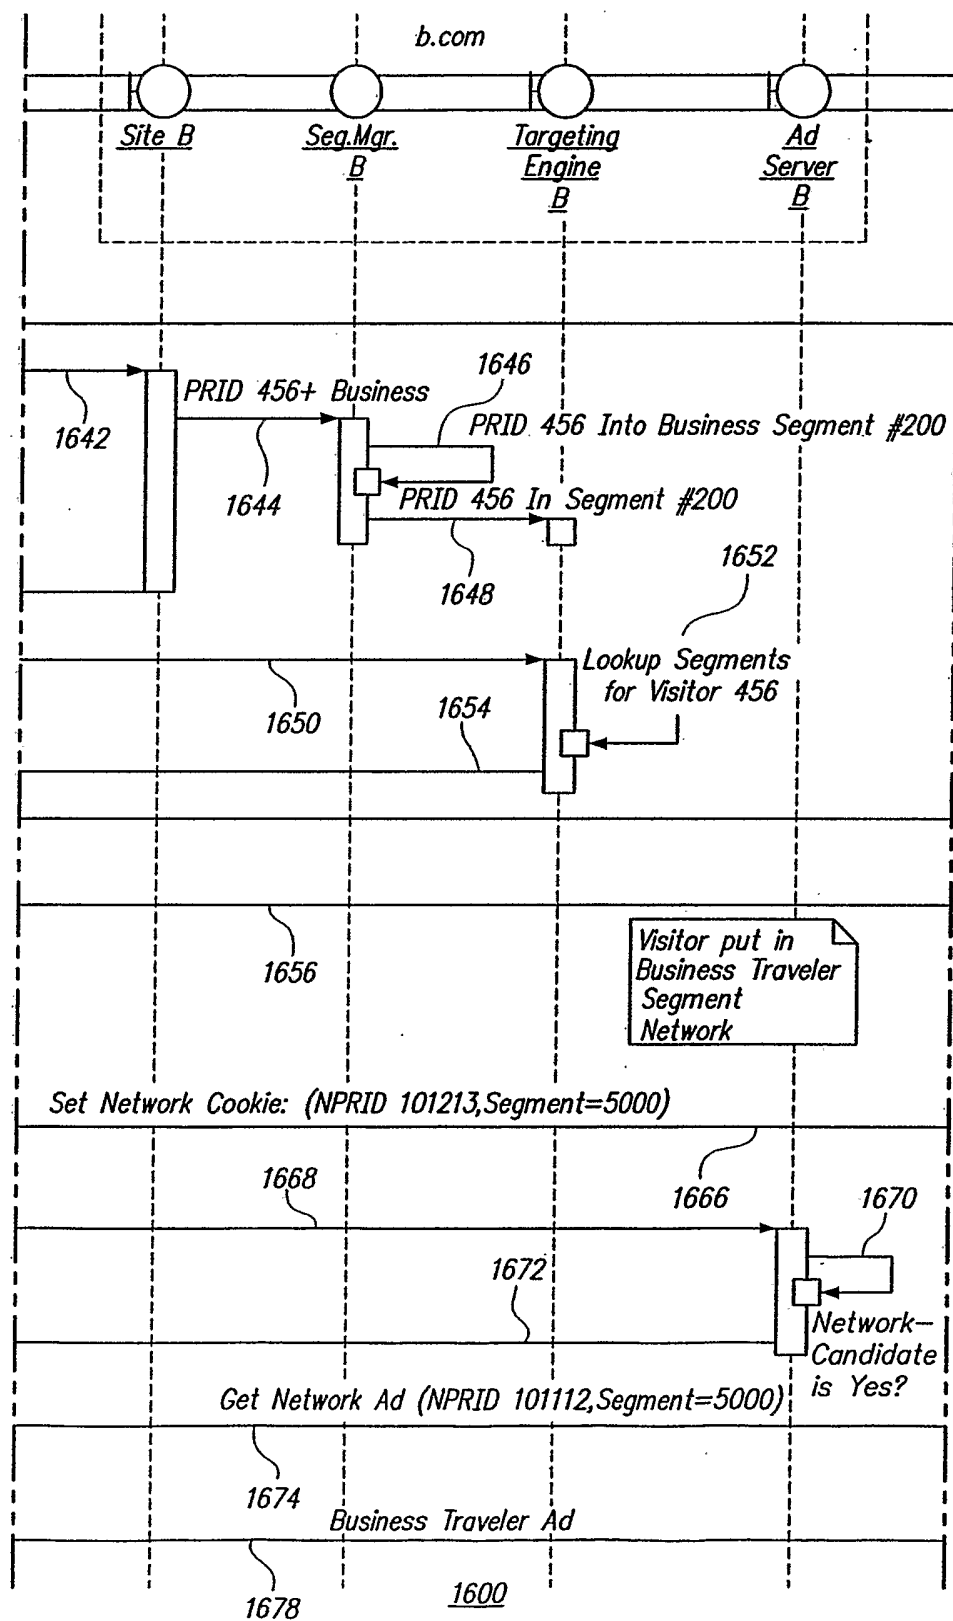


FIG. 16B-3

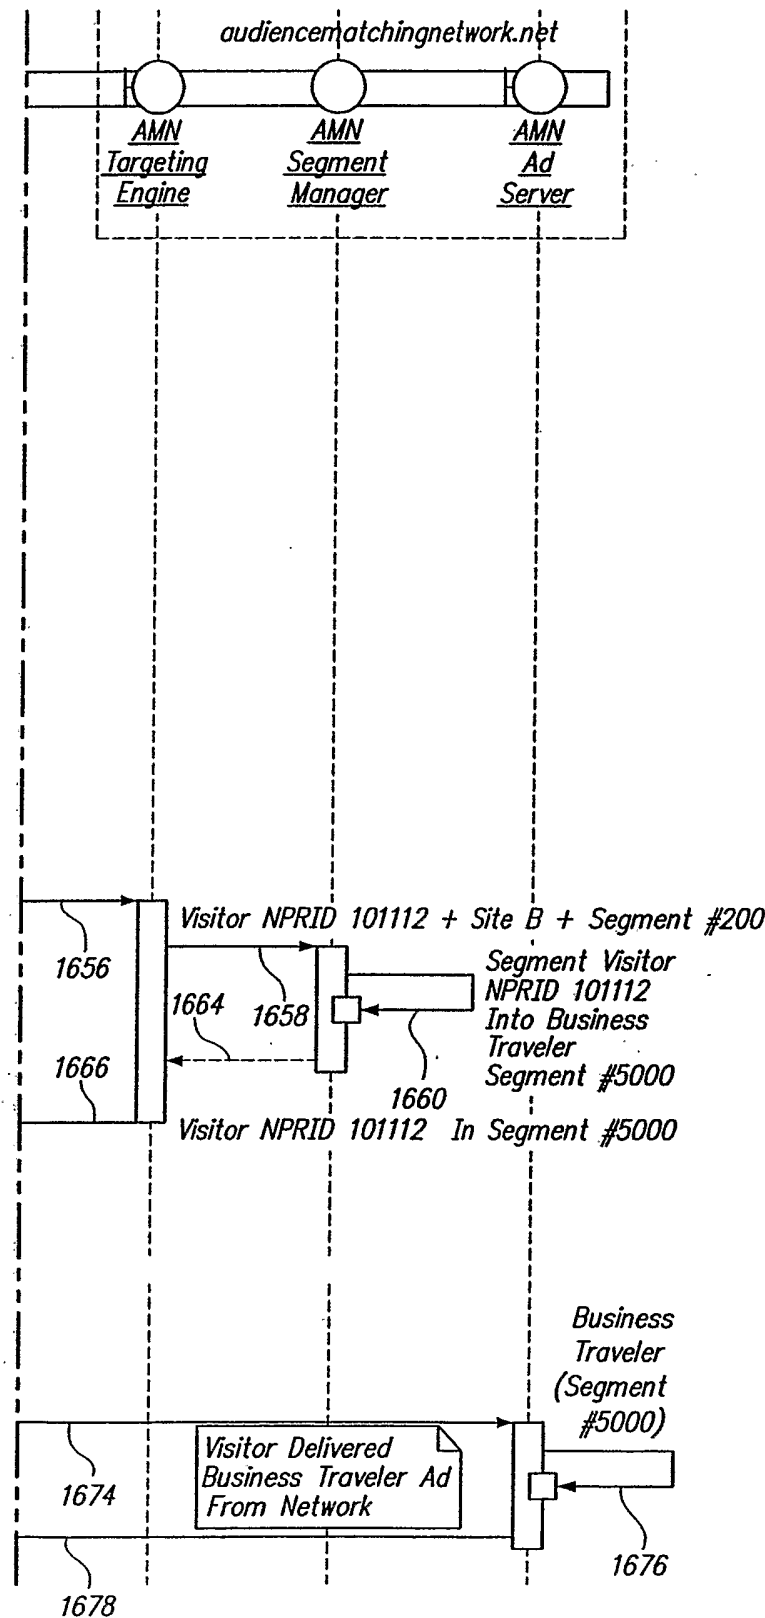


FIG. 16B

FIG. 16B-1	FIG. 16B-2	FIG. 16B-3
---------------	---------------	---------------


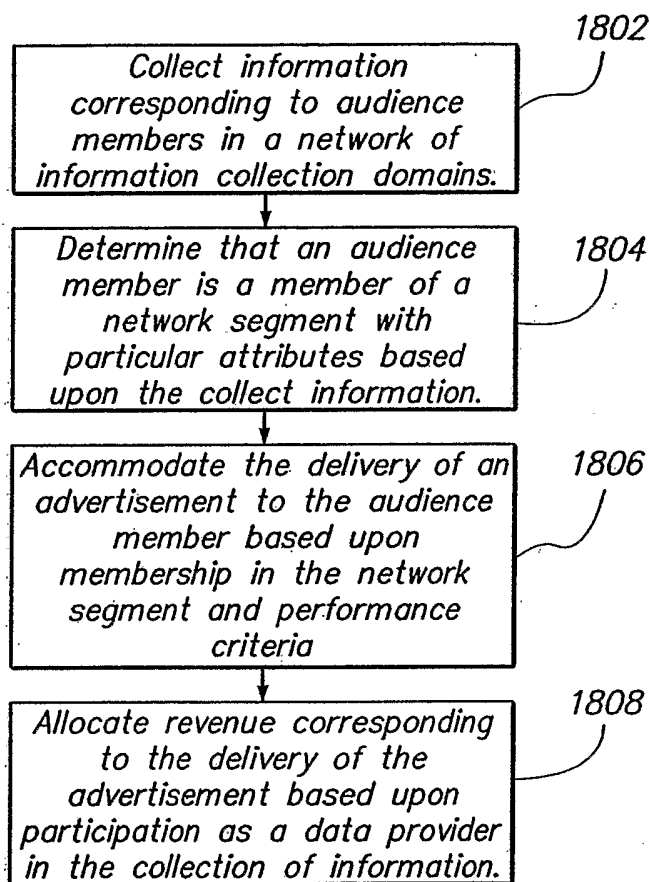
FIG. 17

FIG. 18

1800

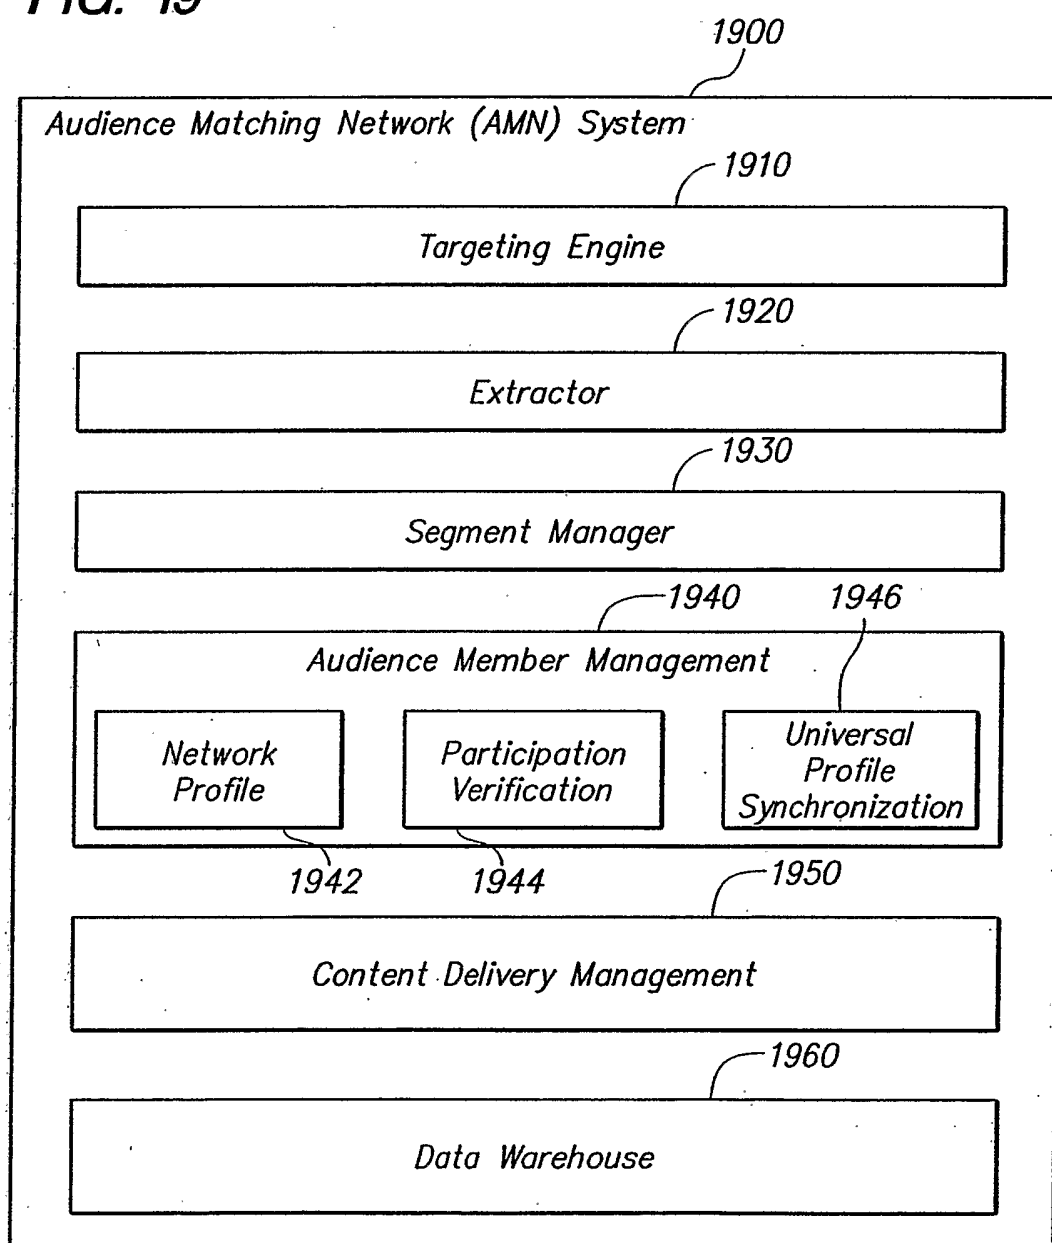
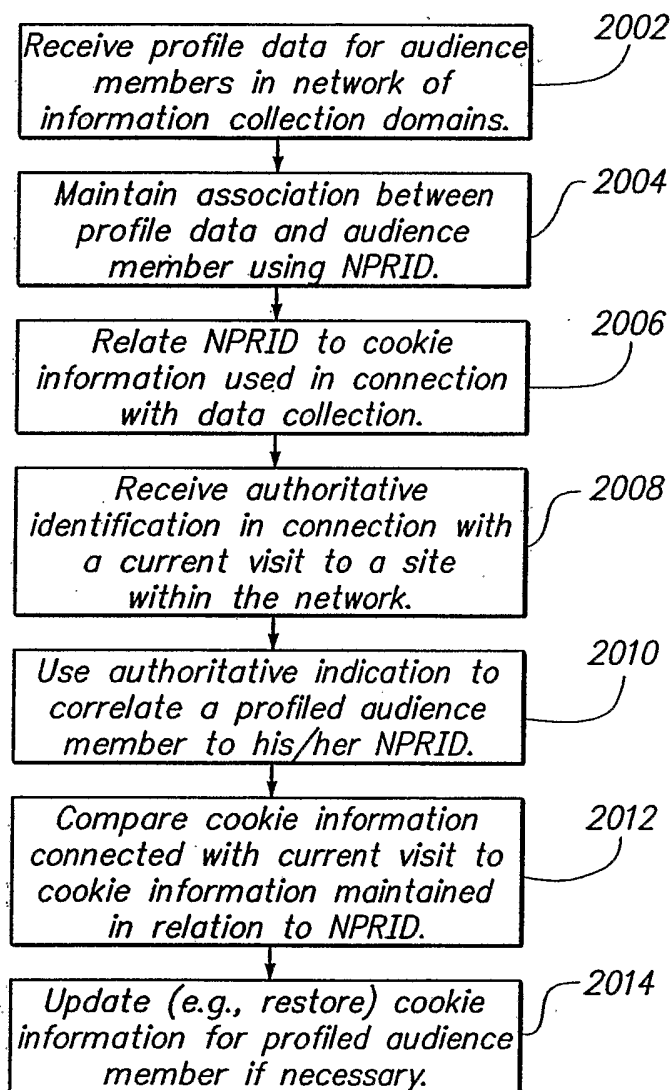
FIG. 19

FIG. 202000