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

Canadian Intellectual Property Office
An agency of Industry Canada

CA 2714057 C 2016/11/08

(11)(21) **2 714 057**

(12) **BREVET CANADIEN**
CANADIAN PATENT

(13) **C**

(86) **Date de dépôt PCT/PCT Filing Date:** 2009/02/03
(87) **Date publication PCT/PCT Publication Date:** 2009/08/13
(45) **Date de délivrance/Issue Date:** 2016/11/08
(85) **Entrée phase nationale/National Entry:** 2010/08/04
(86) **N° demande PCT/PCT Application No.:** US 2009/032976
(87) **N° publication PCT/PCT Publication No.:** 2009/100072
(30) **Priorité/Priority:** 2008/02/04 (US12/025,239)

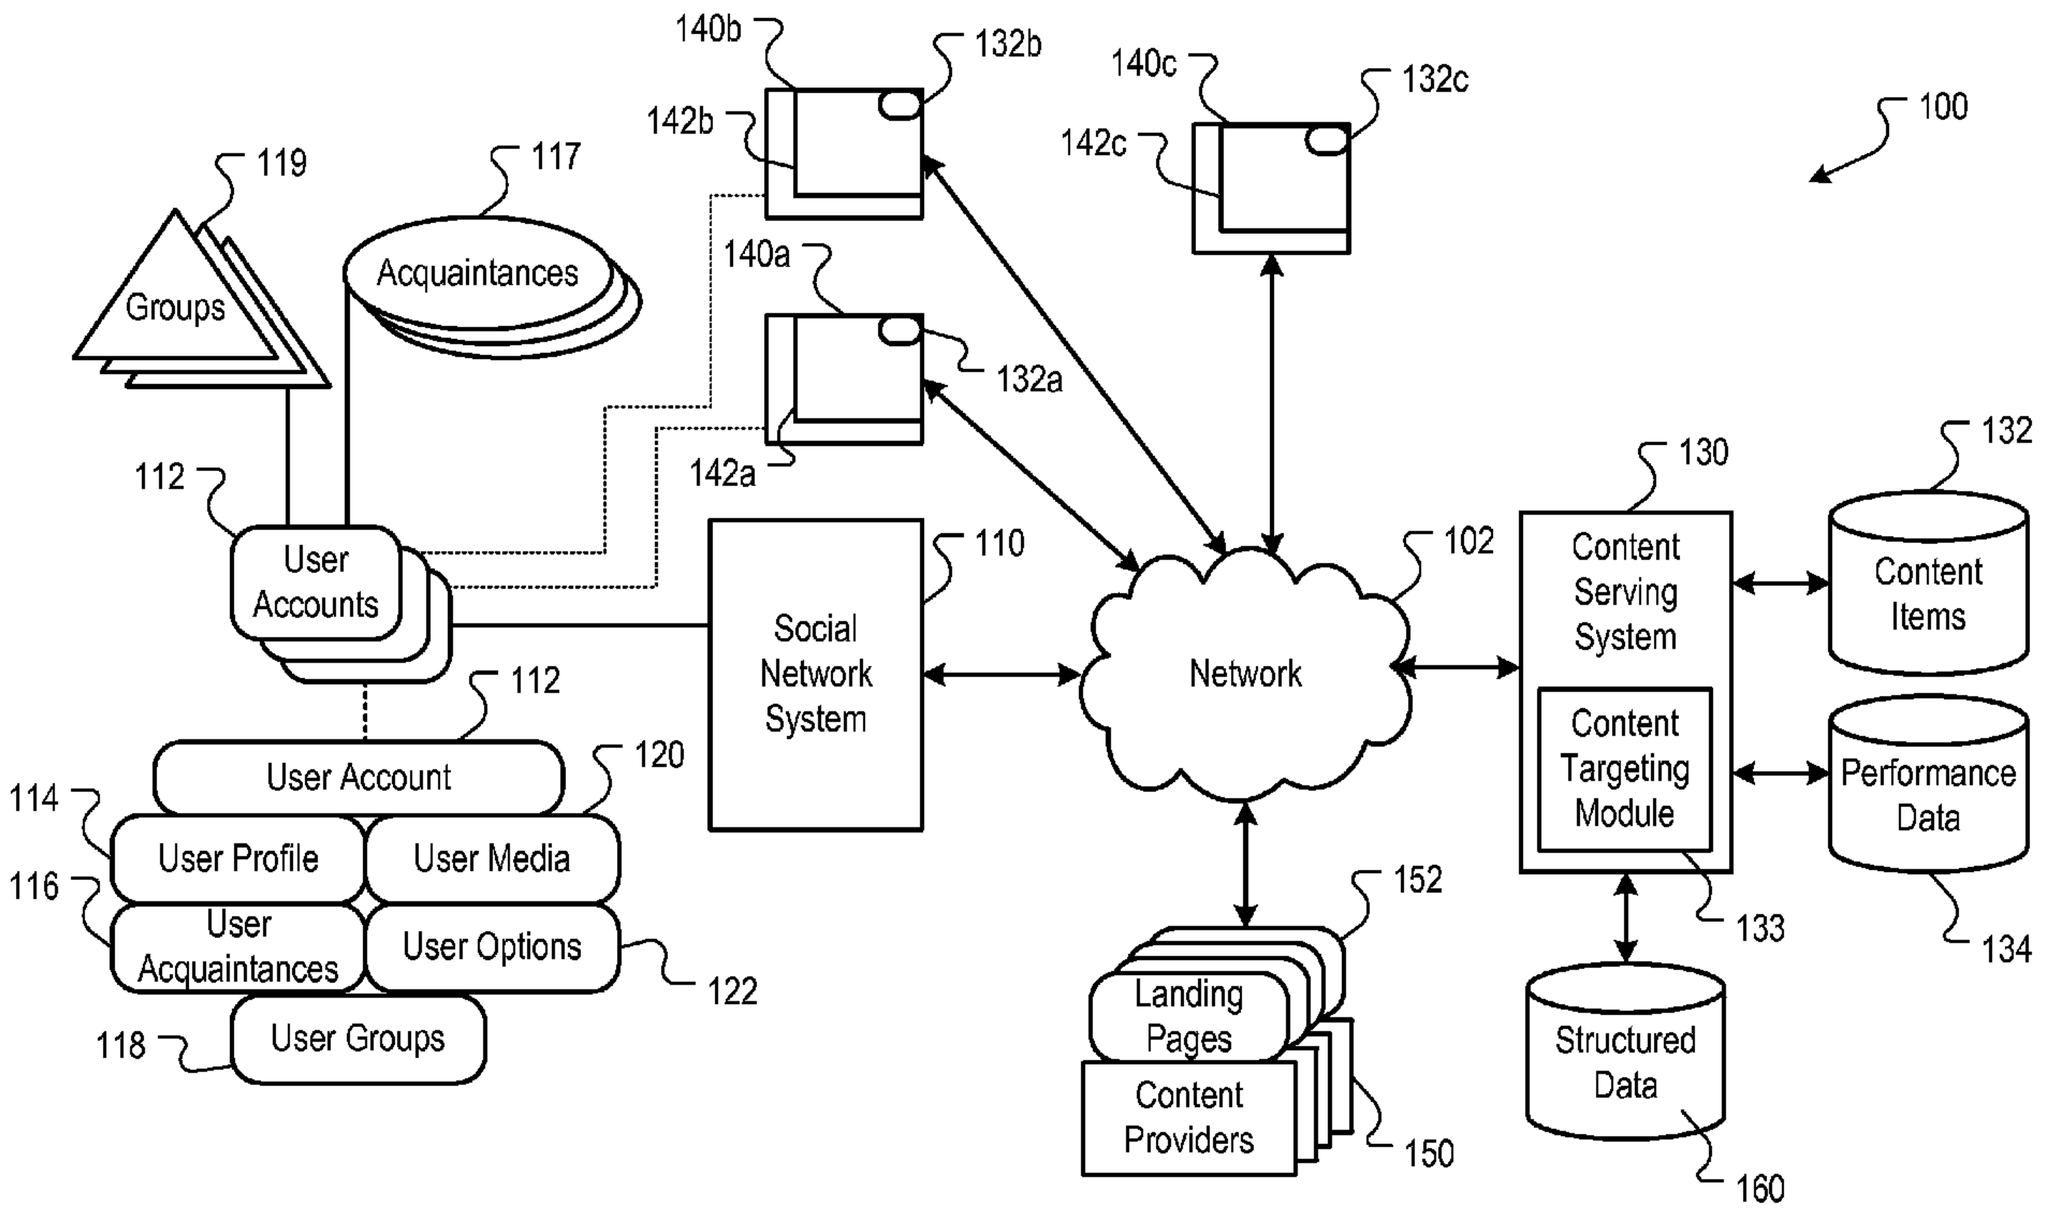
(51) **Cl.Int./Int.Cl.** ***G06F 17/00*** (2006.01),
G06Q 30/02 (2012.01)

(72) **Inventeurs/Inventors:**
DATAR, MAYUR, US;
MILLER, JASON C., US;
HOCHBERG, MIKE, US;
RABII, BAHMAN, US;
NANCE, MEGAN, US;
TUNG, JULIE, US;
HARMSEN, JEREMIAH, US;
...

(73) **Propriétaire/Owner:**
GOOGLE INC., US

(74) **Agent:** SMART & BIGGAR

(54) **Titre : PUBLICITE CIBLANT L'UTILISATEUR**
(54) **Title: USER-TARGETED ADVERTISING**



(57) **Abrégé/Abstract:**
Structured data sets including one or more attributes are identified, each structured data set associated with, for example, a particular user. Values corresponding the at least one of the one or more attributes in each structured data set are identified. A condition established by a content provider, for instance, an advertiser, is compared to the identified values to determine if the condition is satisfied. When the condition is satisfied, one or more content items are identified to the users associated with the structured data sets containing identified values satisfying the condition.

CA 2714057 C 2016/11/08

(11)(21) **2 714 057**

(13) **C**

(72) **Inventeurs(suite)/Inventors(continued):** TUNGUZ-ZAWISLAK, TOMASZ J., US; PEREZ-BERGQUIST, ANDRES S., US

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

(19) World Intellectual Property Organization
International Bureau

(43) International Publication Date
13 August 2009 (13.08.2009)

PCT

(10) International Publication Number
WO 2009/100072 A3

(51) **International Patent Classification**:
G06Q 30/00 (2006.01)

(21) **International Application Number**:
PCT/US2009/032976

(22) **International Filing Date**:
3 February 2009 (03.02.2009)

(25) **Filing Language**: English

(26) **Publication Language**: English

(30) **Priority Data**:
12/025,239 4 February 2008 (04.02.2008) US

(71) **Applicant** *(for all designated States except US)*: **GOOGLE INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, California 94043 (US).

(72) **Inventors; and**
(75) **Inventors/Applicants** *(for US only)*: **DATAR, Mayur** [IN/US]; 131 Saratoga Avenue, #3201, Santa Clara, California 95051 (US). **MILLER, Jason C.** [US/US]; 1860 Washington Street, #304, San Francisco, California 94109 (US). **HOCHBERG, Mike** [US/US]; 1348 Bobolink Circle, Sunnyvale, California 94087 (US). **RABII, Bahman** [US/US]; 2910 King Street, Berkeley, California 94703 (US). **NANCE, Megan** [US/US]; 1273 Lakeside Drive, #3164, Sunnyvale, California 94085 (US). **TUNG, Julie** [US/US]; 100 E. Middlefield Road, #6G, Mountain View, California 94043 (US). **HARMSEN, Jeremiah** [US/US]; 950 High School Way, #3101, Mountain View, California 94041 (US). **TUNGUZ-ZAWISLAK, Tomasz J.** [US/US]; 2055 Sacramento Street, #1007, San Francisco, California 94109 (US). **PEREZ-BERGQUIST, Andres S.** [US/US]; 1901 Rock Street, #109, Mountain View, California 94043 (US).

(74) **Agent**: **BORCHERS, William R.**; Fish & Richardson P.C., P.O. Box 1022, Minneapolis, Minnesota 55440-1022 (US).

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

(84) **Designated States** *(unless otherwise indicated, for every kind of regional protection available)*: ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

[Continued on next page]

(54) **Title**: USER-TARGETED ADVERTISING

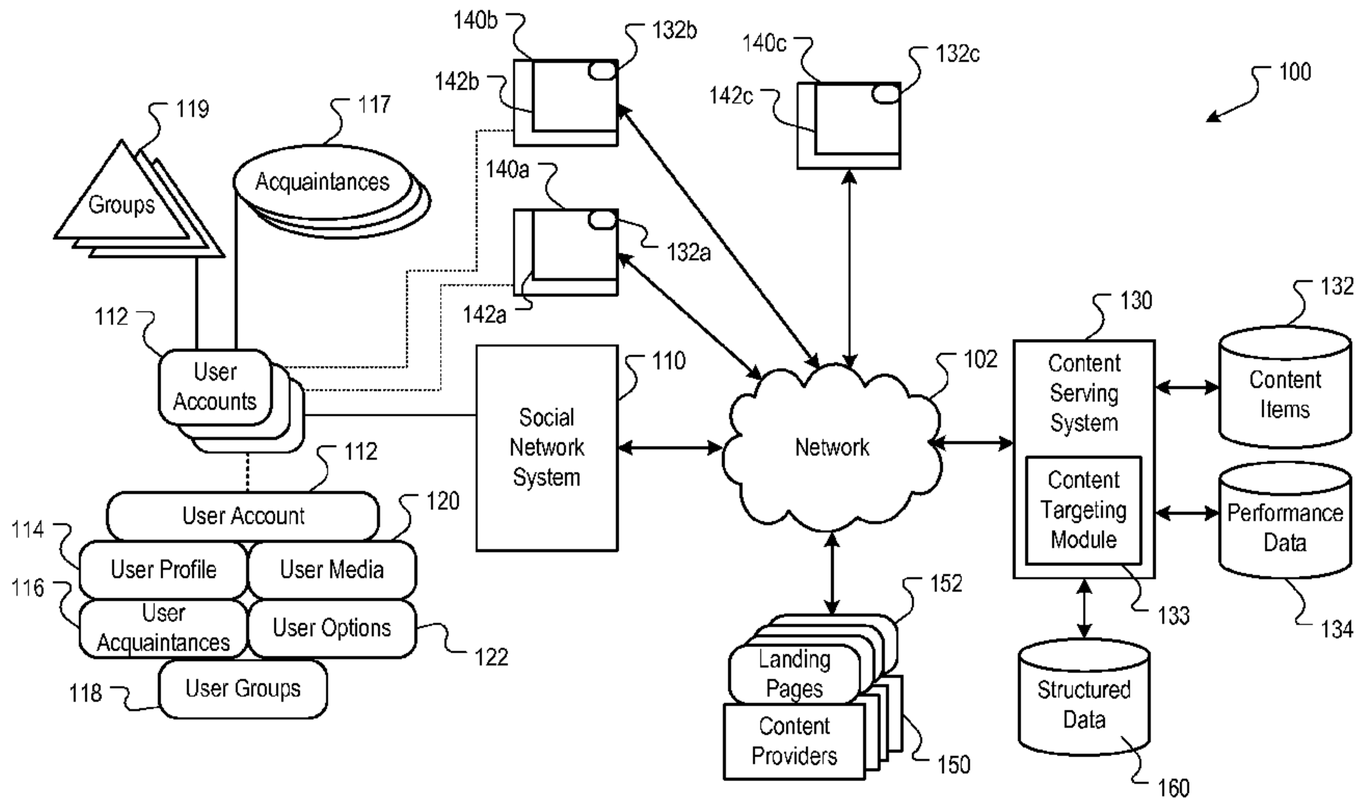


FIG. 1

(57) **Abstract**: Structured data sets including one or more attributes are identified, each structured data set associated with, for example, a particular user. Values corresponding the at least one of the one or more attributes in each structured data set are identified. A condition established by a content provider, for instance, an advertiser, is compared to the identified values to determine if the condition is satisfied. When the condition is satisfied, one or more content items are identified to the users associated with the structured data sets containing identified values satisfying the condition.

WO 2009/100072 A3

WO 2009/100072 A3

— *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*

Published:

— *with international search report (Art. 21(3))*

— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))*

(88) Date of publication of the international search report:
12 November 2009

95569-23

USER-TARGETED ADVERTISING

BACKGROUND

[0001] The subject matter of this application is generally related to information presentation.

[0002] Advertising using traditional media, such as television, radio, newspapers and magazines, is well known. Unfortunately, even when armed with demographic studies and entirely reasonable assumptions about the typical audience of various media outlets, advertisers recognize that a non-trivial portion of their ad budget is simply wasted. Moreover, it is very difficult to identify and eliminate such waste. Recently, advertising over more interactive media has become popular. For example, as the number of people using the Internet has exploded, advertisers have come to appreciate media and services offered over the Internet as a potentially powerful way to advertise.

[0003] Web site-based advertisements ("ads" also referred to as "Web ads") are some times presented to their advertising audience in the form of "banner ads" – i.e., a rectangular box that includes graphic components. When a member of the advertising audience (hereinafter referred to as a "viewer" or "user" without loss of generality) selects one of these banner ads by clicking on it, embedded hypertext links typically direct the viewer to the advertiser's Web site. This process, wherein the viewer selects an ad, is commonly referred to as a "click-through" ("click-through" is intended to cover any user selection.). The ratio of

the number of click-throughs to the number of impressions of the ad (i.e., the number of times an ad is displayed) is commonly referred to as the "click-through rate" of the ad.

[0004] Online social networks have become popular for professional and/or social networking, and are thus popular locations for placing advertisements. Some online social networks provide content items that may be of interest to users, including identification of other users and/or groups that may of interest to a user. Advertisement placement on social networking sites is challenging because such sites have users with many different interests. Accordingly, many advertisements, e.g., advertisements directed to particular products, may not be of interest to many users of an online social network.

SUMMARY

[0005] Systems, methods, and computer program products described in this specification establish a structured data set for each of a number of users. Content providers establish conditions that are compared to values within the structured data sets to identify users that satisfy the conditions. Content providers can display content only to those users whose structured data sets satisfy the conditions. This may be particularly advantageous where a large number of characteristics are known about users, such as users that maintain profiles that are stored and accessible by the system and processes described herein. This can include, for instance, users that are members of online social networks. Additionally, content providers can identify the number of users that satisfy conditions and use that knowledge to further modify the conditions in real-time (i.e., immediately after creating the initial condition) to broaden or restrict the number of users that may receive content. Using an iterative process allows content providers to target a specific audience. This iterative and real-time process may be advantageous, for instance, during the creation of an advertising campaign where content providers can refine their conditions repeatedly in minutes to identify an appropriate subset of users to receive content.

[0006] In general, one aspect of the subject matter described in this specification can be embodied in methods that include the actions of identifying a plurality of structured data sets, each structured data set associated with a user and including one or more attributes, identifying values corresponding to at least one of the one or more attributes in each structured data set, comparing a condition established by a content provider to one or more of the identified values, and determining a number of structured data sets having the one or more of the identified values that satisfy the condition established by the content provider.

[0007] According to a feature, the method includes receiving a revised condition from the content provider. According to another feature, the method can include comparing the revised condition to one or more of the identified values. The method can also include determining a number of structured data sets having the one or more of the identified values that satisfy the revised condition established by the content provider. One or more content items may also be displayed to users associated with the structured data sets having the one or more of the identified values that satisfy the condition established by the content provider.

[0008] According to yet another feature, the method can include identifying a click through rate associated with displayed one or more content items.

[0009] According to another features, the one or more attributes represent characteristics of the user associated with each structured data set. The one or more attributes can include age, gender, number of friends, interests, zip code, or length of time since previous login.

[0010] According to still additional features, the method can include receiving the condition from the content provider, where the condition identifies one or more required values associated with at least one of the one or more attributes. The one or more required values can be compared with the identified values. In another feature, the method can determine when the identified values satisfy the condition. According to yet another feature,

the one or more attributes include one or more public attributes and/or one or more private attributes. The one or more public attributes can be displayed to the content provider.

[0011] Other embodiments of this aspect include corresponding systems, apparatus, and computer program products.

[0011a] In an aspect, there is provided a method performed by one or more processing devices, comprising: identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user accounts; processing, by the one or more processing devices, the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user; storing, in response to the processing, one or more revised attributes and revised values for each user; identifying the revised values corresponding to the one or more revised attributes for each user; comparing a first condition established by a content provider to one or more of the revised values; tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for each user satisfy the first condition; determining a first set of users having revised values that satisfy the first condition; identifying one or more content items for presentation to the first set of users; determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users; determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and presenting the first expected performance and the second expected performance to the content provider.

[0011b] In another aspect, there is provided one or more non-transitory machine-readable storage devices storing instructions that are executable by one or more processing devices to perform operations comprising: identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user

accounts; processing the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user; storing, in response to the processing, one or more revised attributes and revised values for each user; identifying the revised values corresponding to the one or more revised attributes for each user; comparing a first condition established by a content provider to one or more of the revised values; tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for each user satisfy the first condition; determining a first set of users having revised values that satisfy the first condition; identifying one or more content items for presentation to the first set of users; determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users; determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and presenting the first expected performance and the second expected performance to the content provider.

[0011c] In a further aspect, there is provided a system comprising: a data processing apparatus; and a memory couple to the data processing apparatus having instructions stored thereon which, when executed by the data processing apparatus cause the data processing apparatus to perform operations comprising: identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user accounts; processing the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user; storing, in response to the processing, one or more revised attributes and revised values for each user; identifying the revised values corresponding to the one or more revised attributes for each user; comparing a first condition established by a content provider to one or more of the revised values; tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for

each user satisfy the first condition; determining a first set of users having revised values that satisfy the first condition; identifying one or more content items for presentation to the first set of users; determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users; determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and presenting the first expected performance and the second expected performance to the content provider.

[0012] One or more of the following advantages may be achieved. Content providers such as advertisers can target a certain market segment for marketing. Rather than displaying content to an entire user population content providers can restrict content to a certain subset of users that the content providers can identify based on their attributes, i.e., characteristics. As an alternate to restricting ads to a subset of users, content providers may pay more for ads shown to a subset of users.

BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1 is a block diagram of an example system for identifying content items based on structured data sets.

Fig. 2 is a more detailed block diagram of the example system for identifying content items based on structured data sets.

Fig. 3 is a flow diagram of an example process for identifying content items based on structured data sets.

Fig. 4 is a flow diagram of an example process for identifying a number of users whose structured data sets satisfy conditions established by a content provider.

Fig. 5 is a flow diagram of an example process for iteratively permitting a content provider to modify the conditions that determine the users to which content items will be displayed.

Fig. 6 is a block diagram of an example computer system that can be utilized to implement the systems and methods described herein.

DETAILED DESCRIPTION

[0013] Fig. 1 is a block diagram of an example system 100 for identifying content items, such as advertisements, based on structured data sets. In some implementations, each structured data set corresponds to a particular user. Alternatively, a structured data set can correspond to all members of a class of users. Structured data sets include structured data that, in some implementations, includes a number of attribute-value pairs. Attributes may be common to many users and identify characteristics of users (e.g., age, gender, weight, etc.), whereas the value corresponding to each attribute in each structured data set defines a characteristic of that user. For instance, a structured data set for an individual may include the following attribute-values pairs:

Age : 35 yrs
Gender : Male
Number of Friends : 72
Interests : Cars, Reading
Zip code: 30067
Days Since Last Login: 10

where 'Age', 'Gender', 'Number of Friends', 'Interests', 'Zip code', and 'Days Since Last Login' are attributes, and '35 yrs', 'Male', '72', 'Cars', 'Reading', '30067', and '10' are values. The attributes and values associated with each are used to determine whether each user satisfies conditions established by a content provider, as explained in greater detail below.

[0014] In some implementations, structured data may be identified from data collected and/or stored by a social network system 110 that can, for example, host numerous user accounts 112. An example social network system can include Orkut, hosted by Google, Inc., of Mountain View, CA. Other social networks can, for example, include school alumni websites, an internal company web site, dating networks, or the like. Each user account 112

can, for example, include user profile data 114, user acquaintances data 116, user groups data 118, user media data 120, and user options data 122. Other user data can also be included in the user accounts 112. The data in each user account can be used to populate the structured data 160 for each user, and more particularly, the attribute-value pairs. Structure data 160 may be stored internal or external (as shown in Fig. 1) to the social network system 110 or may be stored at a location accessible anywhere on a network 102, such as the Internet. In some cases, where the social network is different from the content server, a copy of the structured data may be provided by the social network to the content server and maintained separately and updated regularly.

[0015] The user profile data 114 can, for example, include general demographic data about a user, such as age, sex, home location, interests, etc. In some implementations, the user profile data 114 can also include professional information, e.g., occupation, educational background, etc., and other data, such as contact information. In some implementations, the user profile data 114 can include open profile data, e.g., free-form text that may be typed by the user into fields for various subjects, e.g., "Job Description," "Favorite Foods," "Hobbies," etc., and constrained profile data, e.g., binary profile data selected by check boxes, radio buttons, etc., or predefined selectable profile data, e.g., income ranges, zip codes, etc. In some implementations, some or all or the user profile data 114 can be classified as public or private profile data, e.g., data that can be shared publicly or data that can be selectively or otherwise not shared. According to some implementations, profile data 114 not classified as private data can, for example, be classified as public data, e.g., data that can be viewed by any user accessing the social network system 110.

[0016] The user acquaintances data 116 can, for example, define user acquaintances 117 associated with a user account 112. In some implementations, user acquaintances 117 can include, for example, users associated with other user accounts 112 that are classified as "friends," e.g., user accounts 112 referenced in a "friends" or "buddies" list. Other

acquaintances 117 can also be defined, e.g., professional acquaintances, client acquaintances, family acquaintances, etc. In some implementations, the user acquaintance data 116 for each user account 112 can, for example, be specified by users associated with each user account 112, and thus can be unique for each user account 112.

[0017] The user groups data 118 can, for example, define user groups 119 to which a user account 112 is associated. In an implementation, user groups 119 can, for example, define an interest or topic, e.g., "Wine," "Open Source Chess Programming," "Travel Hints and Tips," etc. In some implementations, the user groups 119 can, for example, be categorized, e.g., a first set of user groups 119 can belong to an "Activities" category, a second set of user groups 119 can belong to an "Alumni & Schools" category, etc.

[0018] The user media data 120 can, for example, include one or more content items, such as user documents, e.g., web pages. A document can, for example, include a file, a combination of files, one or more files with embedded links to other files, etc. The files can be of any type, such as text, audio, image, video, hyper-text mark-up language documents, etc. In the context of the Internet, a common content item is a Web page. Other content items can also be defined by the user media data 120.

[0019] The user options data 122 can, for example, include data specifying user options, such as e-mail settings, acquaintance notification settings, chat settings, password and security settings, etc. Other option data can also be included in the user options data 122.

[0020] In some implementations, the content serving system 130 can directly, or indirectly, enter, maintain, identify, and track content items 132. The content items 132 can, for example, include a web page or other content document, or text, graphics, video, audio, mixed media, etc. In some implementations, the content items 132 are advertisements. The advertisements can, for example, be in the form of graphical ads, such as banner ads, text only ads, image ads, audio ads, video ads, ads combining one of more of any of such

components, etc. The advertisements can also include embedded information, such as links, meta-information, and/or machine executable instructions.

[0021] In some implementations, user devices 140a, 140b and 140c can communicate with the social network system 110 over the network 102, such as the Internet. The user devices 140 can be any device capable of receiving the user media data 120, such as personal computers, mobile devices, cell phones, personal digital assistants (PDAs), television systems, etc. The user devices 140 can be associated with user accounts 112, e.g., the users of user devices 140a and 140b can be logged-in members of the social network system 110, having corresponding user accounts 112a and 112b. Additionally, the user devices 140 may not be associated with a user account 112, e.g., the user of the user device 142c may not be a member of the social network system 110 or may be a member of the social network system 110 that has not logged in.

[0022] In some implementations, upon a user device 140 communicating a request for media data 120 of a user account 112, or a request for some other social network media data to the social network system 110, the social network system 110 can, for example, provide the user media data 120 or the other requested media data to user device 140. In some implementations, the user media data 120 or other requested media data can include an embedded request code, such as one or more Javascript code snippets. In other implementations, the social network system 110 can insert the embedded request code into the user media data 120 or the other requested media data when served to a user device 140.

[0023] The user device 140 can render the requested data, e.g., a user web page, or some other social network 110 page, in a presentation environment 142, e.g., in a web browser application. Upon rendering the requested data, the user device 140 executes the request code, which causes the user device 140 to issue a content request, e.g., an advertisement request, to the content serving system 130. In response, the content serving system 130 can provide one or more content items 132 to the user device 140. For example,

the content items 132a, 132b and 132c can be provided to the user devices 140a, 140b and 140c, respectively. In some implementations, the content items 132a, 132b and 132c are presented in the presentation environments 142a, 142b and 142c, respectively.

[0024] Rather than displaying content to an entire user population content providers 150 can restrict content to a certain subset of users that the content providers 150 can identify based on their attributes, i.e., user characteristics. According to some implementations, the content serving system 130 includes a content targeting module 133 that can identify users that should receive content items 132 based on structured data sets 160 for those users, such as users associated with the user accounts 112.

[0025] The content targeting module 133 receives conditions from content providers 150 and compares those conditions to values within each user's structured data set to identify those users whose attributes and values satisfy the content provider conditions. As an example, a content provider 150 may establish the condition that users receive content, or a particular type of content, only when those users' structured data sets meet the following attribute-value conditions: "Age: between 25 and 35", "Zip code: 95051", and "Number of Friends > 10." This permits the content provider 150 to target certain subsets of users.

[0026] Although displayed as part of the content serving system 130 in the implementation shown in Fig. 1, the content targeting module 133 can be integrated into the social network system 110. In other implementations, the content targeting module 133 can be a separate system in data communication with the social network system 110 and/or the content server system 130. The content targeting module 133 can be implemented in software and executed on a processing device, such as the computer system 600 of Fig. 6. Example software implementations include C, C++, Java, or any other high-level programming language that may be utilized to produce source code that can be compiled into executable instructions. Other software implementations can also be used, such as applets, or

interpreted implementations, such as scripts, etc. Alternatively, the content targeting module 133 can be implemented in hardware or a combination of hardware and software.

[0027] Content providers 150 can use the identified number of users to further modify the conditions to broaden or restrict the number of users that satisfy the conditions. Using an iterative process allows content providers 150 to target a specific audience. This may be particularly advantageous where a large number of characteristics are known about users, such as users that maintain profiles and data stored within a social networking system 110.

[0028] The iterative process of modifying conditions to identify a desired subset of users can occur in real-time based on current structured data and performance information. For instance, a content provider such as an advertiser may adjust the conditions using a graphical user interface (GUI) at the content targeting module 133. An advertiser can repeatedly refine conditions in minutes to identify an appropriate subset of users to receive content. Each time a condition is established it is compared by the content targeting module 133 against currently stored structured data and performance data to provide the content provider with real-time results. Content providers can view in real time what how many users will receive impressions, the anticipated click through rate of those impressions, and the like, for each content-provider entered condition.

[0029] According to some implementations, content providers may be charged or be willing to pay a higher content placement fee (e.g., advertising fee) for placement of content directed to a specific set of users that satisfy the content provider's conditions. Thus, once a content provider has established conditions that the content provider feels captures the correct target audience, the content provider may be charged a premium for content placement for those target users. Although the advertiser may pay more for content placement in front of the target group, the advertiser may be willing to pay a greater amount for placement of content, e.g., advertisements, that are provided to a group of users that are likely to provide the advertiser with a high return on the advertiser's investment.

[0030] According to an implementation, conditions may be compared by the content targeting module 133 to structured data sets during periods of low traffic or at other times regardless of whether a content provider has requested identification of the users that may satisfy the conditions. When this occurs, each user's structured data set may be tagged by the content targeting module 133 as satisfying (or failing to satisfy) one or more conditions and/or combination of conditions. This optimization reduces the need for the content targeting module 133 to constantly compare all values within structured data sets to conditions in real-time or in near-real time.

[0031] According to some implementations, the content providers 150 can, for example, include web sites having "landing pages" 152 that a user is directed to when the user clicks (e.g., interacts with) an advertisement 132 presented on page provided from the social networking system 110. For example, the content providers 150 can provide content items 132 in the form of "creatives," which are advertisements that may include text, graphics and/or audio associated with the advertised service or product, and a link to a web site.

[0032] In some implementations, the content serving system 130 can monitor and/or evaluate performance data 134 related to the content items 132. For example, the performance of each advertisement 132 can be evaluated based on a performance metric, such as a click-through rate, a conversion rate, or some other performance metric. A click-through can occur, for example, when a user of a user device, e.g., user device 140a, selects or "clicks" on an advertisement, e.g. the advertisement 132a. The click-through rate can be a performance metric that is obtained by dividing the number of users that clicked on the advertisement or a link associated with the advertisement by the number of times the advertisement was delivered. For example, if advertisement is delivered 100 times, and three persons clicked on the advertisement, then the click-through rate for that advertisement is 3%.

[0033] A "conversion" occurs when a user, for example, consummates a transaction related to a previously served advertisement. What constitutes a conversion may vary from case to case and can be determined in a variety of ways. For example, a conversion may occur when a user of the user device 140a clicks on an advertisement 132a, is referred to the advertiser's Web page, such as one of the landing pages 152, and consummates a purchase before leaving that Web page. Other conversion types can also be used. A conversion rate can, for example, be defined as the ratio of the number of conversions to the number of impressions of the advertisement (i.e., the number of times an advertisement is rendered) or the ratio of the number of conversions to the number of selections. Other types of conversion rates can also be used.

[0034] Other performance metrics can also be used. The performance metrics can, for example, be revenue related or non-revenue related. In some implementations, the performance metrics can be parsed according to time, e.g., the performance of a particular content item 132 may be determined to be very high on weekends, moderate on weekday evenings, but very low on weekday mornings and afternoons, for example. Performance data can be provided to the content providers 150 to assist the content provider in delivering, modifying, and determining content, such as advertisements, that is delivered to users. Performance data can also assist the content providers 150 in modifying the conditions that identify users that receive content.

[0035] Fig. 2 is a more detailed block diagram of the example system 100 for identifying content items based on structure data sets. In the example system 100 of Fig. 2, the content targeting module 133 can utilize attribute-value pairs from structured data sets to determine those users that satisfy conditions established by one or more content providers 150.

[0036] In some implementations, attributes and their corresponding values are extracted from user profile data 114a – 114n, user media data 120a-120n, user acquaintances

data 116a-116n, user groups data 118a-118n, and user options data 122a-122n of user accounts 112a – 112n. Attributes can include characteristics such as age, gender, zip code, and the like. Attributes can also include traits such as personality, values, attitudes, interests, lifestyles, etc. of users associated with the user accounts 112a – 112n. Other types of attributes are possible.

[0037] According to some implementations, information provided by users in free-form text fields may be extracted by other system components, such as a natural language processor (not illustrated), to identify attributes and/or values. For instance, a user may identify their favorite hobbies by filling out a text field under the caption “Hobbies” by typing the sentence "I really enjoy hiking, especially long hikes when you can camp out for a few days. Indoor activities don't interest me at all, and I really don't like boring outdoor activities like gardening.” A natural language processor can identify the attribute as “Hobbies” and values as “Hiking”. As can be seen from the previous example, data associated with a user may identify things a user is disinterested in. Attribute value pairs can therefore include this information as well. Thus, in the above example, the attribute “Hobbies” may include a value “Gardening” that is modified or also paired with a negative indicator to indicate that the user is disinterested in a particular subject. A content provider 150 can establish a condition, e.g., that identifies users that fall within certain specifications, including disinterest in a particular subject.

[0038] In some implementations, an example process flow for identifying structured data sets begins when the content targeting module 133 receives the user profile data 114a – 114n, user media data 120a-120n, user acquaintances data 116a-116n, user groups data 118a-118n, and user options data 122a-122n of user accounts 112a – 112. The content serving system 130 and/or content targeting module 133 can, for example, be configured to extract attribute-value pairs from the received data. According to some implementations, the content

targeting module 133 can utilize another component, such as a natural language processor or the like, to extract attribute-value pairs.

[0039] In some implementations, data 114, 116, 118, 120, 122 associated with a particular user account, e.g. user account 112a can, for example, be provided to the content serving system 130 by a user device 140 associated with the particular user account 112, e.g., the user device 140a upon which a logged in session for the particular user account 112a has been established. For example, if the user device 140a renders a web page from the social network system 110 that includes a content item request instruction, a corresponding content item request 126a that includes the user profile data 114a – 114n, user media data 120a-120n, user acquaintances data 116a-116n, user groups data 118a-118n, and/or user options data 122a-122n can be provided to the content serving system 130.

[0040] Other schemes can also be used to provide the structured data 160 to the content server system 130. For example, the content serving system 130 may partner with the social network system 110 to receive user account 112 identifiers in corresponding content requests 126a – 126n. In response to receiving the user account 112 identifiers, the content serving system 130 can transmit the user account 112 identifiers to the social network system 110 in a request for corresponding structured data 160. The social network system 110 can, in response, provide the corresponding structured data 160 to the content serving system 130. As mentioned earlier, in some implementations the social network system 110 periodically provides an updated copy of the structured data sets to the content serving system 130.

[0041] According to some implementations, the content serving system 130 and/or content targeting module 133 receives one or more conditions from the content providers 150 and compares the conditions to structured data corresponding to each user (also referred to as a 'structured data set'). The content targeting module 133 identifies the structured data sets, and thus users, that satisfy the content provider conditions. The content targeting module 133

can identify the users and/or number of users to the content provider 150, which can then optionally revise the conditions to target a broader or narrower group of users. The subset of users that satisfy the conditions may be provided to content providers in real time. When conditions are revised the content targeting module 133 re-compares the structured data sets with the revised conditions to identify users that satisfy the revised conditions. The content serving system 130 can then serve content provided from the content providers 150 to users that satisfy the conditions.

[0042] According to some implementations, structured data 160 can be provided to advertisers, e.g., the content providers 150. According to some implementations, only a portion of the structure data may be provided to the content providers 150. Therefore, each attribute-value pair may be identified as a public or private, and the content providers 150 may only receive public attribute-value pairs. According to other implementations, the content providers may be provided only with some or all attributes, and no values. Alternatively, the content providers may be provided only with some or all values, and no attributes.

[0043] Providing this information to the content providers 150 can, for example, facilitate the targeting of content items 132, e.g., advertisements, to social networking sites that have users having many different interests. For example, the structured data 160 for the social network system 110 may identify a set of highly relevant marketing verticals for which many of the users of the social network system 110 may have a strong interest. Accordingly, one or more content providers 150 can target advertisements related to the verticals to users of the social network system 110.

[0044] In some implementations, the content targeting module 133 can process the user accounts 112a – 112n periodically, e.g., monthly, to revise the structured data 160. Other processing triggers, e.g., changes in the user account 112 corpus, can also be used. In some implementations, the content targeting module 133 can process the user accounts 112a

– 112n in an offline batch process. The processing results can, for example, be stored and accessed during the serving of content. In other implementations, the content targeting module 133 can process the user accounts 112a – 112n in an online process, e.g., in response to a user device 140 submitting a content request to the social network system 110.

[0045] Fig. 3 is a flow diagram of an example process 300 for identifying content items based on structured data sets. The process 300 can, for example, be implemented in the content targeting module 133. In some implementations, the content targeting module 133 may be within the content serving system 130. Other implementations, however, can also be used. For instance, the content targeting module 133 can be integrated into the social network system 110 or in a separate system in data communication with the social network system 110 and/or the content server system 130.

[0046] Stage 301 builds structured data sets using data associated with each user profile, including, for instance, user profile data 114a – 114n, user media data 120a-120n, user acquaintances data 116a-116n, user groups data 118a-118n, and/or user options data 122a-122n. For example, the content targeting module 133 can identify and extract attribute and values from each of set of data.

[0047] Stage 302 identifies values in the structured data sets that correspond to attributes. For example, the content targeting module 133 can identify and extract attribute and values from each of set of data. The identified values are compared to content provider conditions in stage 304. For example, the content targeting module 133 can execute the comparison of values in structured data sets to the required values within conditions established by one or more content providers.

[0048] For instance, if a content provider conditions require: “Age: between 25 and 35", and “Zip code: 95051", then the content targeting module 133 can compare the values from the structured data sets that include the attributes “Age” and ‘Zip code” and determine whether the values satisfy the required values. The content targeting module 133 can store

the results of the comparison such that users whose structured data sets satisfy the conditions may be targeted by content providers to receive content, e.g., an advertisement.

[0049] Stage 306 identifies one or more content items for display to users whose structured data sets include values satisfying the conditions. For example, the content serving system 130 can identify content items 132 that the content provider wishes to present to users.

[0050] Other processes for identifying content items based on structured data sets can also be used.

[0051] Fig. 4 is a flow diagram of an example process 400 for identifying a number of users whose structured data sets satisfy conditions established by a content provider.

[0052] The process 400 can, for example, be implemented in the content targeting module 133. In some implementations, the content targeting module 133 may be within the content serving system 130. Other implementations, however, can also be used. For instance, the content targeting module 133 can be integrated into the social network system 110 or in a separate system in data communication with the social network system 110 and/or the content server system 130.

[0053] Stage 401 builds structured data sets using data associated with each user profile, including, for instance, user profile data 114a – 114n, user media data 120a-120n, user acquaintances data 116a-116n, user groups data 118a-118n, and/or user options data 122a-122n. For example, the content targeting module 133 can identify and extract attribute and values from each of set of data.

[0054] Stage 402 identifies values in the structured data sets that correspond to attributes. For example, the content targeting module 133 can identify and extract attribute and values from each of set of data. The identified values are compared to content provider conditions in stage 404. For example, the content targeting module 133 can execute the comparison of values in structured data sets to the required values within conditions

established by one or more content providers. The content targeting module 133 can store the results of the comparison such that users whose structured data sets satisfy the conditions may be targeted by content providers to receive content, e.g., an advertisement.

[0055] Stage 406 identifies the number of users whose structured data sets include values satisfying the conditions. For example, the content serving system 130 can identify the number of users that may be served content items 132 under the content provider's current conditions. According to some implementations, this information may be provided to the content provider, which may use the information to determine whether the conditions should be revised to target a greater or lesser number of users.

[0056] Other processes for identifying a number of users satisfying content provider conditions can also be used.

[0057] Fig. 5 is a flow diagram of an example process 500 for iteratively permitting a content provider to modify the conditions that determine the users to which content items will be displayed.

[0058] The process 500 can, for example, be implemented in the content targeting module 133. In some implementations, the content targeting module 133 may be within the content serving system 130. Other implementations, however, can also be used. For instance, the content targeting module 133 can be integrated into the social network system 110 or in a separate system in data communication with the social network system 110 and/or the content server system 130.

[0059] Stage 501 receives content provider conditions. The conditions may be transmitted, for instance, from a content provider to the content serving system 130 or may be entered by a content provider using the content serving system 130 and/or content targeting module 133. For example, an advertiser may utilize a GUI at the content targeting module 133 to enter the conditions. This may be facilitated, for instance, using a series of pull down menus in which attributes may be selected, along with Boolean operators (e.g., 'and', 'or',

etc.). Values may also be entered by content providers in a similar manner, such as manually or via pull down menus that provide content providers with the ability to select ranges, minimum values, maximum values, and the like.

[0060] Stage 502 compares the conditions received in Stage 501 with identified values in structured data sets corresponding to a group of users. For example, the content targeting module 133 can execute the comparison of values in structured data sets to the required values within conditions established by one or more content providers.

[0061] Stage 504 determines the number of users whose structured data sets include values satisfying the content provider conditions. This process may be implemented, for example, the content targeting module 133. This information is provided to the content provider in Stage 506 so that the content provider can use the information to determine whether the conditions should be revised to target a greater or lesser number of users.

[0062] In stage 508, if the content provider determines that the number of users is appropriate, in stage 510 the content targeting module 133 can identify one or more content items to be displayed to the users whose structured data sets include values satisfying the condition. Otherwise, the content provider can revise the conditions to target a greater or lesser number of users, and the process repeats starting at stage 501.

[0063] Other processes for iteratively permitting a content provider with information to determine whether conditions should be adjusted may also be used. For instance, instead of providing content providers with the total number of users that satisfy conditions, the content providers may be provided with metrics, such as the number of conversion or click-throughs that are anticipated by presenting content to a subset of users identified by the conditions.

[0064] Fig. 6 is a block diagram of an example computer system 600 that can be utilized to implement the systems and methods described herein. The system 600 includes a processor 610, a memory 620, a storage device 630, and an input/output device 640. Each of

the components 610, 620, 630, and 640 can, for example, be interconnected using a system bus 650. The processor 610 is capable of processing instructions for execution within the system 600. In one implementation, the processor 610 is a single-threaded processor. In another implementation, the processor 610 is a multi-threaded processor. The processor 610 is capable of processing instructions stored in the memory 620 or on the storage device 630.

[0065] The memory 620 stores information within the system 600. In one implementation, the memory 620 is a computer-readable medium. In one implementation, the memory 620 is a volatile memory unit. In another implementation, the memory 620 is a non-volatile memory unit.

[0066] The storage device 630 is capable of providing mass storage for the system 600. In one implementation, the storage device 630 is a computer-readable medium. In various different implementations, the storage device 630 can, for example, include a hard disk device, an optical disk device, or some other large capacity storage device.

[0067] The input/output device 640 provides input/output operations for the system 600. In one implementation, the input/output device 640 can include one or more of a network interface devices, e.g., an Ethernet card, a serial communication device, e.g., and RS-232 port, and/or a wireless interface device, e.g., and 802.11 card. In another implementation, the input/output device can include driver devices configured to receive input data and send output data to other input/output devices, e.g., keyboard, printer and display devices 660.

[0068] The apparatus, methods, flow diagrams, and structure block diagrams described in this patent document may be implemented in computer processing systems including program code comprising program instructions that are executable by the computer processing system. Other implementations may also be used. Additionally, the flow diagrams and structure block diagrams described in this patent document, which describe particular methods and/or corresponding acts in support of steps and corresponding functions

in support of disclosed structural means, may also be utilized to implement corresponding software structures and algorithms, and equivalents thereof.

[0069] This written description sets forth the best mode of the invention and provides examples to describe the invention and to enable a person of ordinary skill in the art to make and use the invention. This written description does not limit the invention to the precise terms set forth. Thus, while the invention has been described in detail with reference to the examples set forth above, those of ordinary skill in the art may effect alterations, modifications and variations to the examples without departing from the scope of the invention.

CLAIMS:

1. A method performed by one or more processing devices, comprising:

identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user accounts;

processing, by the one or more processing devices, the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user;

storing, in response to the processing, one or more revised attributes and revised values for each user;

identifying the revised values corresponding to the one or more revised attributes for each user;

comparing a first condition established by a content provider to one or more of the revised values;

tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for each user satisfy the first condition;

determining a first set of users having revised values that satisfy the first condition;

identifying one or more content items for presentation to the first set of users;

determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users;

determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and

presenting the first expected performance and the second expected performance to the content provider.

2. The method of claim 1, further comprising receiving a revised condition from the content provider.

3. The method of claim 2, further comprising comparing the revised condition to one or more of the revised values.

4. The method of claim 3, further comprising determining a second set of the users having revised attributes that satisfy the revised condition.

5. The method of claim 1, further comprising displaying the one or more content items to the users in the first set of users after presenting the first expected performance and the second expected performance to the content provider.

6. The method of claim 5, further comprising identifying a click through rate associated with the one or more content items.

7. The method of claim 1, wherein the one or more attributes are selected from a group of attributes comprising age, gender, number of friends, interests, zip code, or length of time since previous login.

8. The method of claim 1, further comprising receiving the first condition from the content provider, where the first condition identifies one or more required values associated with at least one of the one or more revised attributes.

9. The method of claim 8, further comprising comparing the one or more required values with the revised values.

10. The method of claim 1, wherein at least one of the one or more revised attributes is identified as public.

11. The method of claim 10, wherein at least one of the one or more revised attributes is identified as private.

12. The method of claim 11, further comprising determining to display the revised attributes identified as public to the content provider and determining not to display the revised attributes identified as private to the content provider.

13. The method of claim 1, wherein the first expected performance is for a first time of day and the second expected performance is for a second time of day different from the first time of day.

14. The method of claim 1, wherein the first expected performance is for a first day of the week and the second expected performance is for a second day of the week different from the first day of the week.

15. The method of claim 1, wherein determining a first expected performance and determining a second expected performance are performed in substantially real time.

16. The method of claim 1, wherein the first period of time preceded a current time and the second period of time is subsequent to the current time.

17. The method of claim 1, wherein the triggers comprise at least one of changes in the user accounts or predetermined schedules for revising the attributes and values.

18. The method of claim 1, wherein the first expected performance and the second expected performance are presented to the content provider for performing tasks comprising one or more of modifying the first condition or modifying the one or more content items for presentation to the first set of users.

19. One or more non-transitory machine-readable storage devices storing instructions that are executable by one or more processing devices to perform operations comprising:

identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user accounts;

processing the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user;

storing, in response to the processing, one or more revised attributes and revised values for each user;

identifying the revised values corresponding to the one or more revised attributes for each user;

comparing a first condition established by a content provider to one or more of the revised values;

tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for each user satisfy the first condition;

determining a first set of users having revised values that satisfy the first condition;

identifying one or more content items for presentation to the first set of users;

determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users;

determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and

presenting the first expected performance and the second expected performance to the content provider.

20. The one or more non-transitory machine-readable storage devices of claim 19, the operations further comprising:

receiving a revised condition from the content provider.

21. The one or more non-transitory machine-readable storage devices of claim 20, the operations further comprising:

comparing the revised condition to one or more of the revised values.

22. The one or more non-transitory machine-readable storage devices of claim 21, the operations further comprising:

determining a second set of the users having revised attributes that satisfy the revised condition.

23. The one or more non-transitory machine-readable storage devices of claim 19, the operations further comprising:

displaying the one or more content items to the users in the first set of users after presenting the first expected performance and the second expected performance to the content provider.

24. The one or more non-transitory machine-readable storage devices of claim 23, the operations further comprising:

identifying a click through rate associated with the one or more content items.

25. The one or more non-transitory machine-readable storage devices of claim 19, wherein the first expected performance is for a first time of day and the second expected performance is for a second time of day different from the first time of day.

26. The one or more non-transitory machine-readable storage devices of claim 19, wherein the first expected performance is for a first day of the week and the second expected performance is for a second day of the week different from the first day of the week.

27. A system comprising:

a data processing apparatus; and

a memory couple to the data processing apparatus having instructions stored thereon which, when executed by the data processing apparatus cause the data processing apparatus to perform operations comprising:

identifying users, each user being associated with one or more attributes, the one or more attributes comprising user profile data from user accounts;

processing the user accounts to revise attributes and values corresponding to the one or more attributes for each user, the processing being based on one or more triggers, wherein the values comprise user profile data of each user;

storing, in response to the processing, one or more revised attributes and revised values for each user;

identifying the revised values corresponding to the one or more revised attributes for each user;

comparing a first condition established by a content provider to one or more of the revised values;

tagging the revised attributes and revised values for each user, the tagging indicating that the revised values for each user satisfy the first condition;

determining a first set of users having revised values that satisfy the first condition;

identifying one or more content items for presentation to the first set of users;

determining a first expected performance of the one or more content items if the one or more content items are presented to the first set of users, the first expected performance being associated with a first period of time and being based on the first set of users and the revised attributes currently associated with the first set of users;

determining a second expected performance of the one or more content items if the one or more content items are presented to the first set of users, the second expected performance being associated with a second period of time, different from the first

period of time, and being based on the first set of users and the revised attributes currently associated with the first set of users; and

presenting the first expected performance and the second expected performance to the content provider.

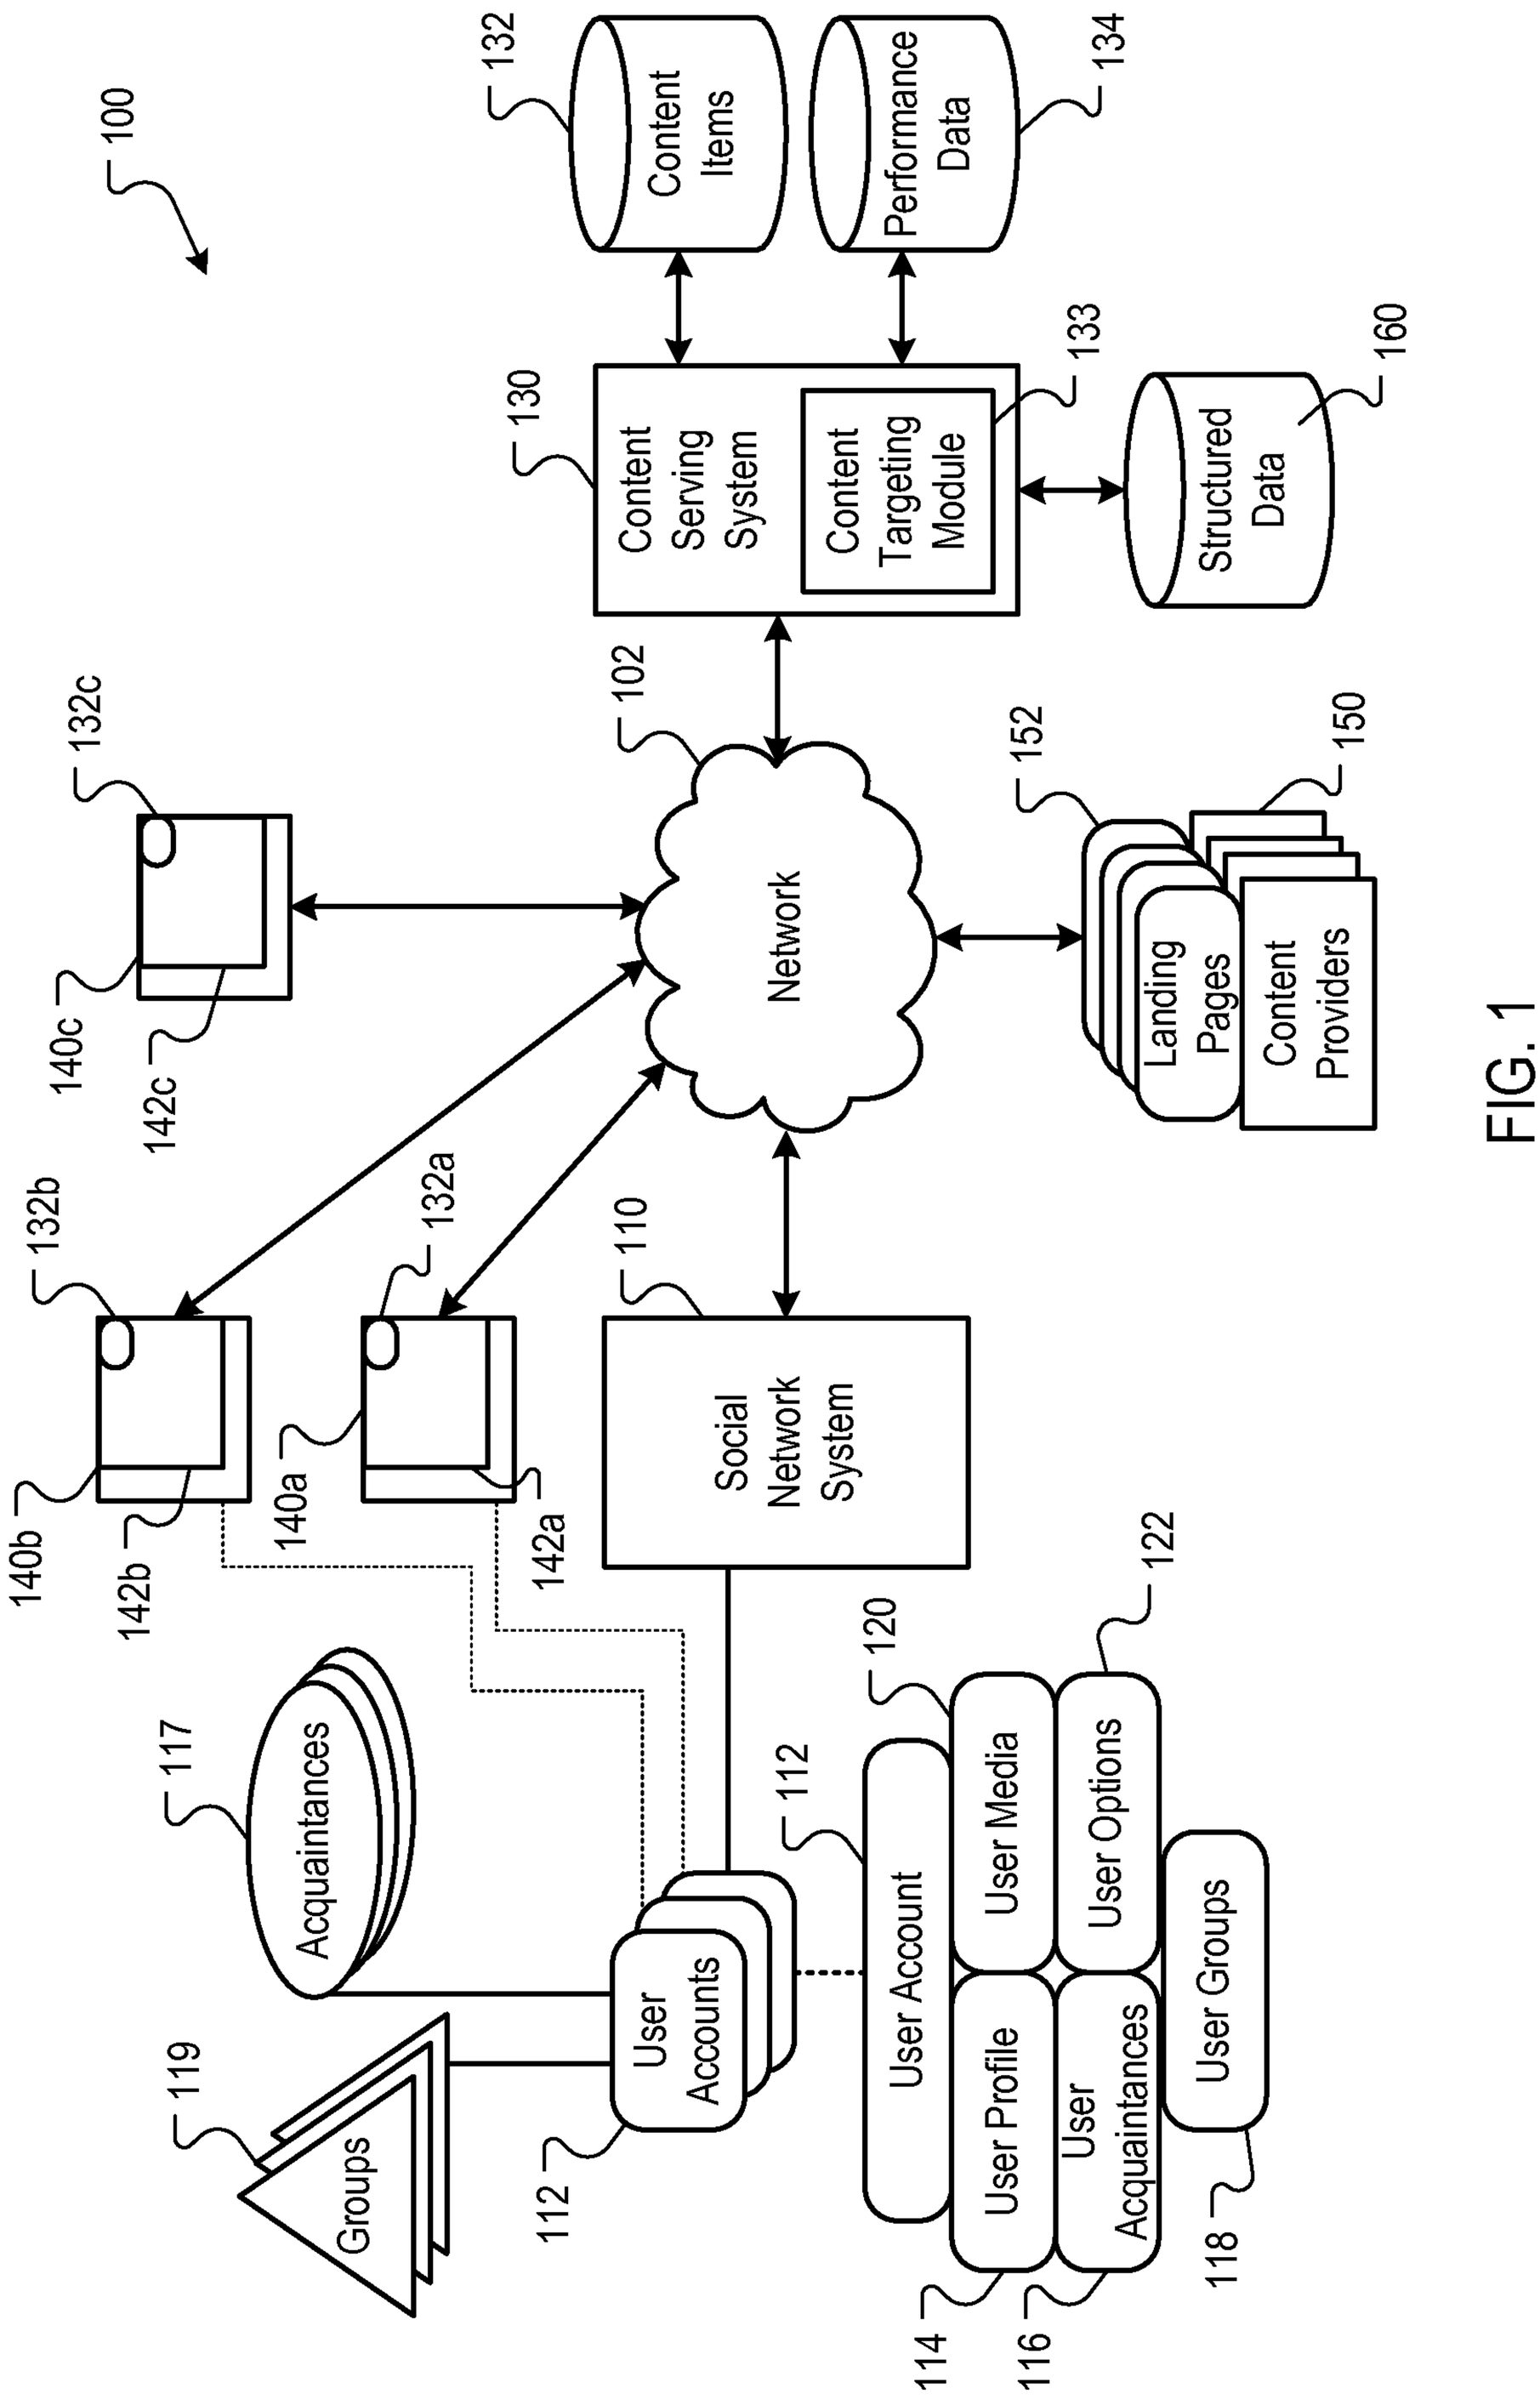

100
132
Content Items
134
Performance Data
130
Content Serving System
Content Targeting Module
133
160
Structured Data
102
Network
132c
140c
142c
152
150
Landing Pages
Content Providers
132b
140b
142b
132a
140a
142a
110
Social Network System
120
122
117
Acquaintances
119
Groups
112
User Accounts
User Account
User Media
User Options
User Profile
User Acquaintances
User Groups
114
116
118

FIG. 1

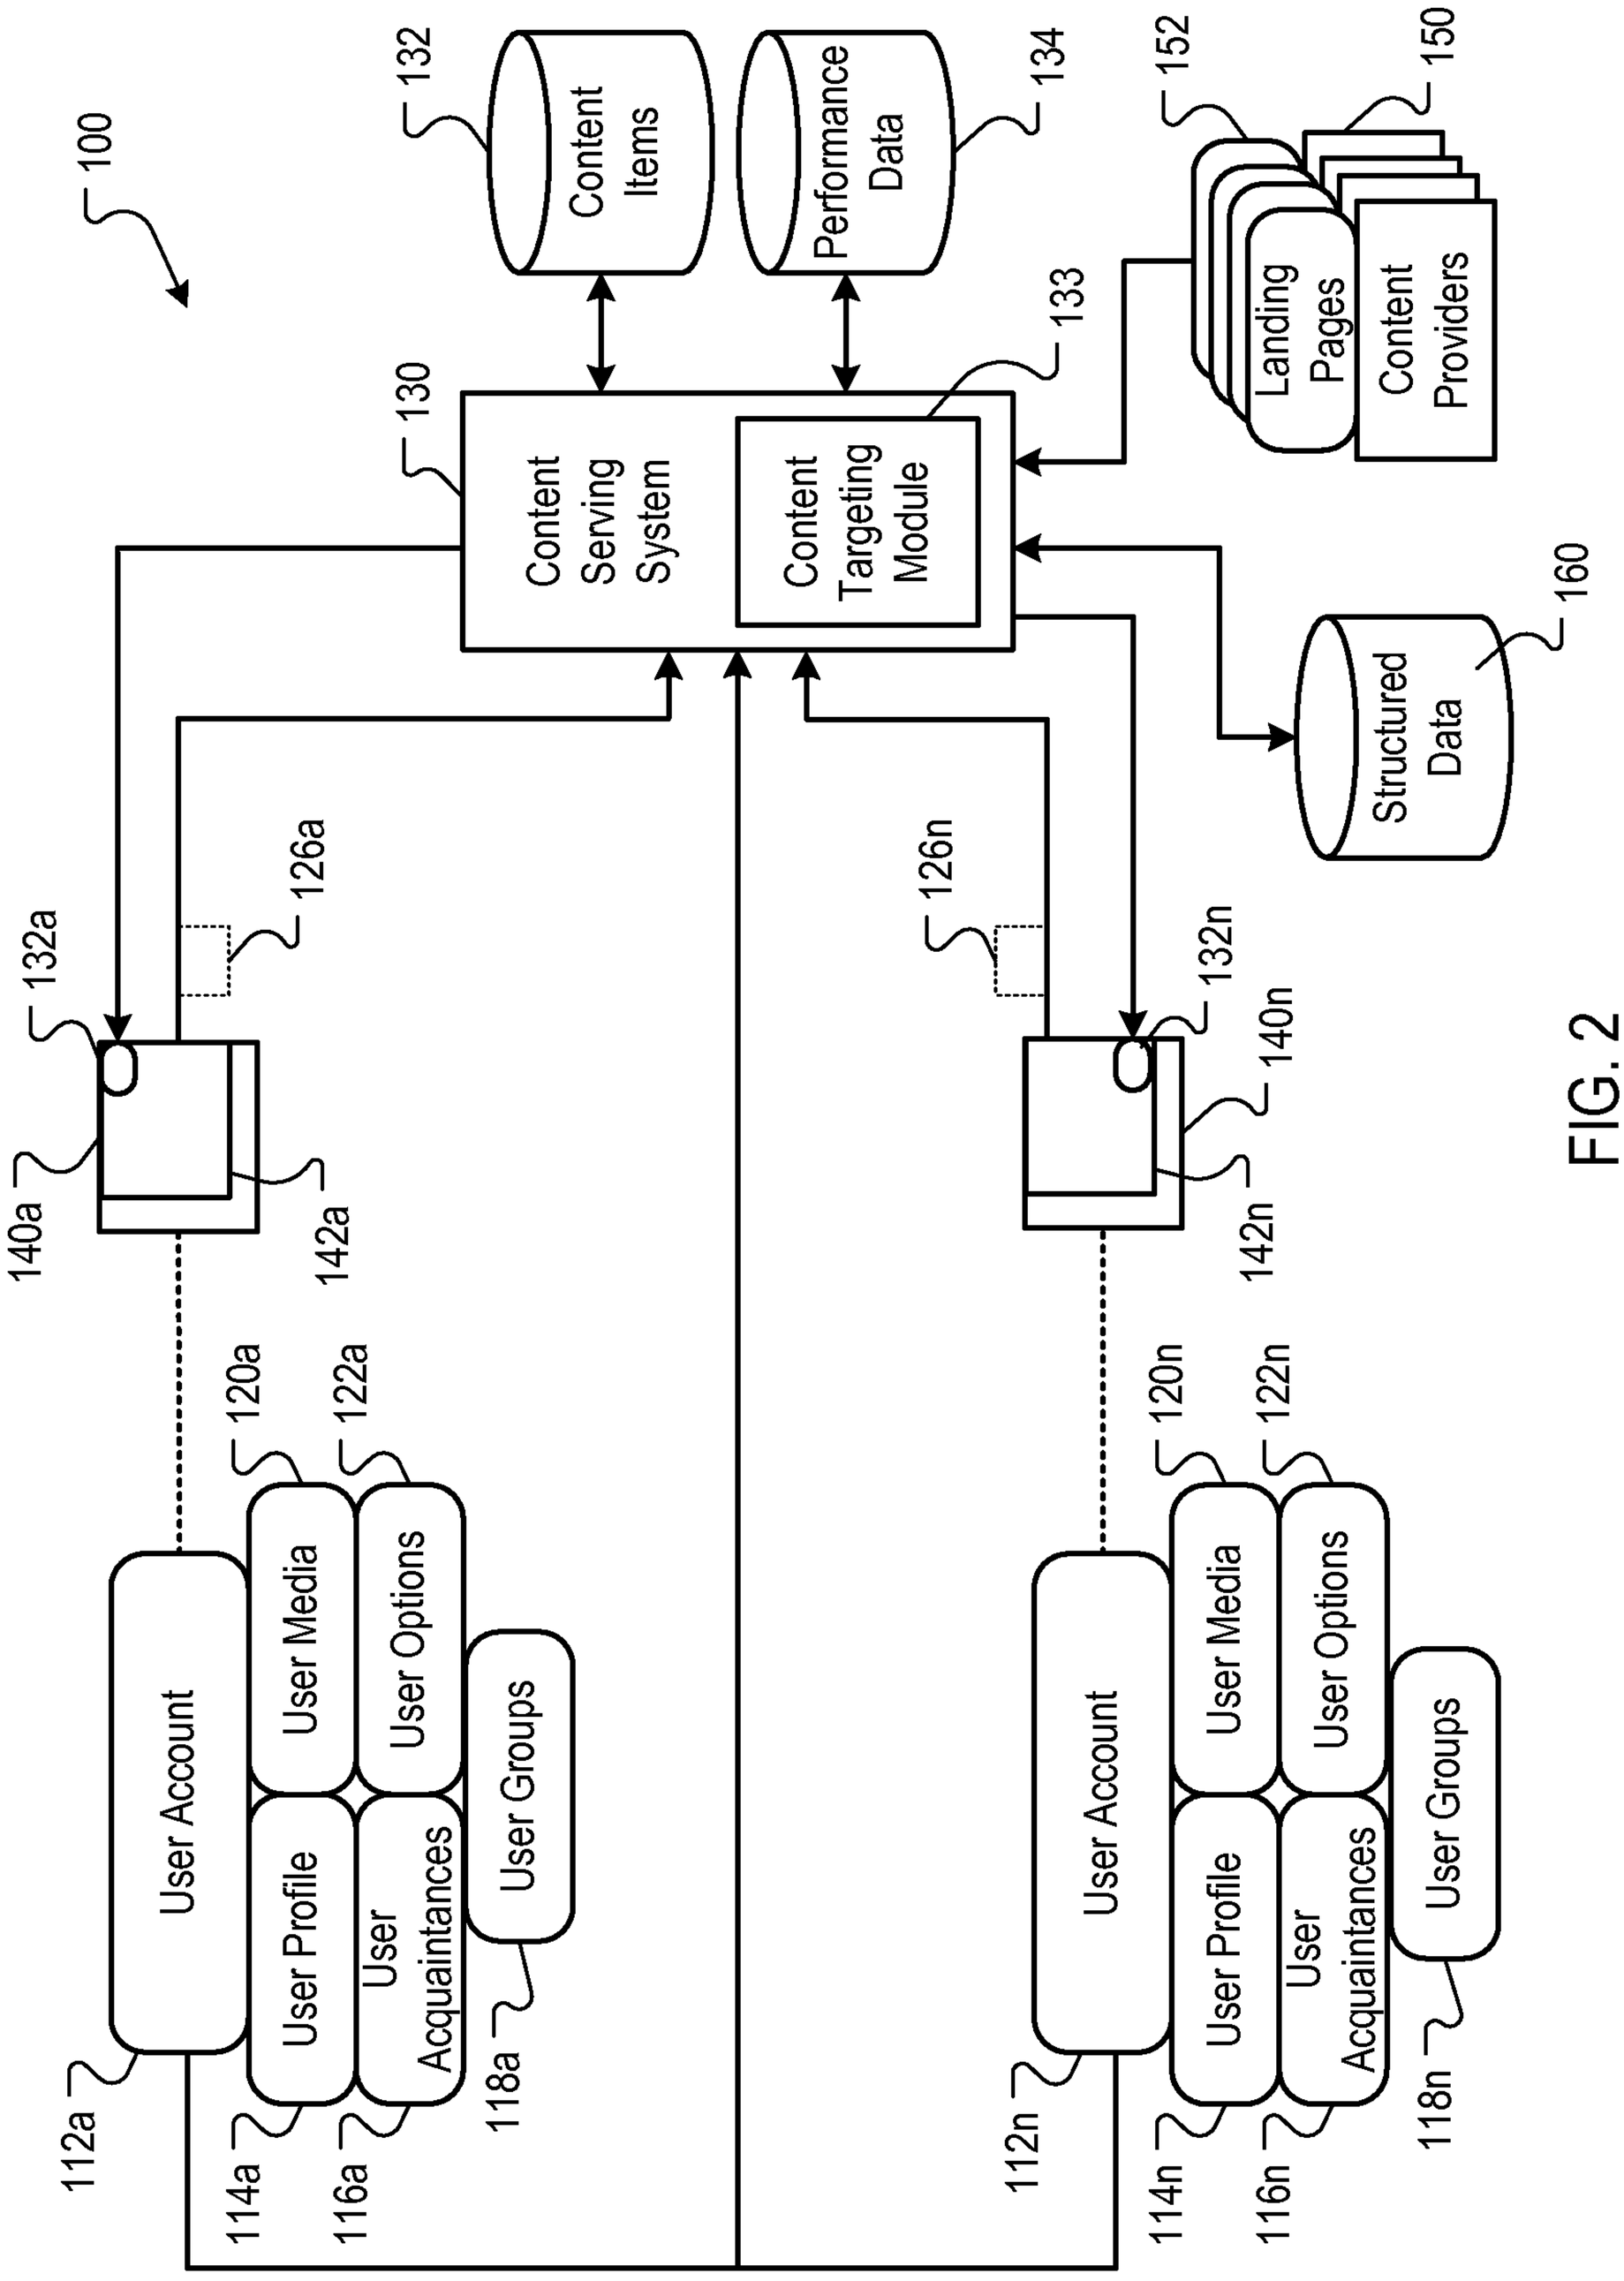

100
Content Items
132
Performance Data
134
152
150
Landing Pages
Content Providers
133
130
Content Serving System
Content Targeting Module
160
Structured Data
126a
132a
140a
142a
120a
122a
User Account
User Media
User Options
User Groups
User Profile
User Acquaintances
112a
114a
116a
118a
126n
132n
140n
142n
120n
122n
User Account
User Media
User Options
User Groups
User Profile
User Acquaintances
112n
114n
116n
118n

FIG. 2

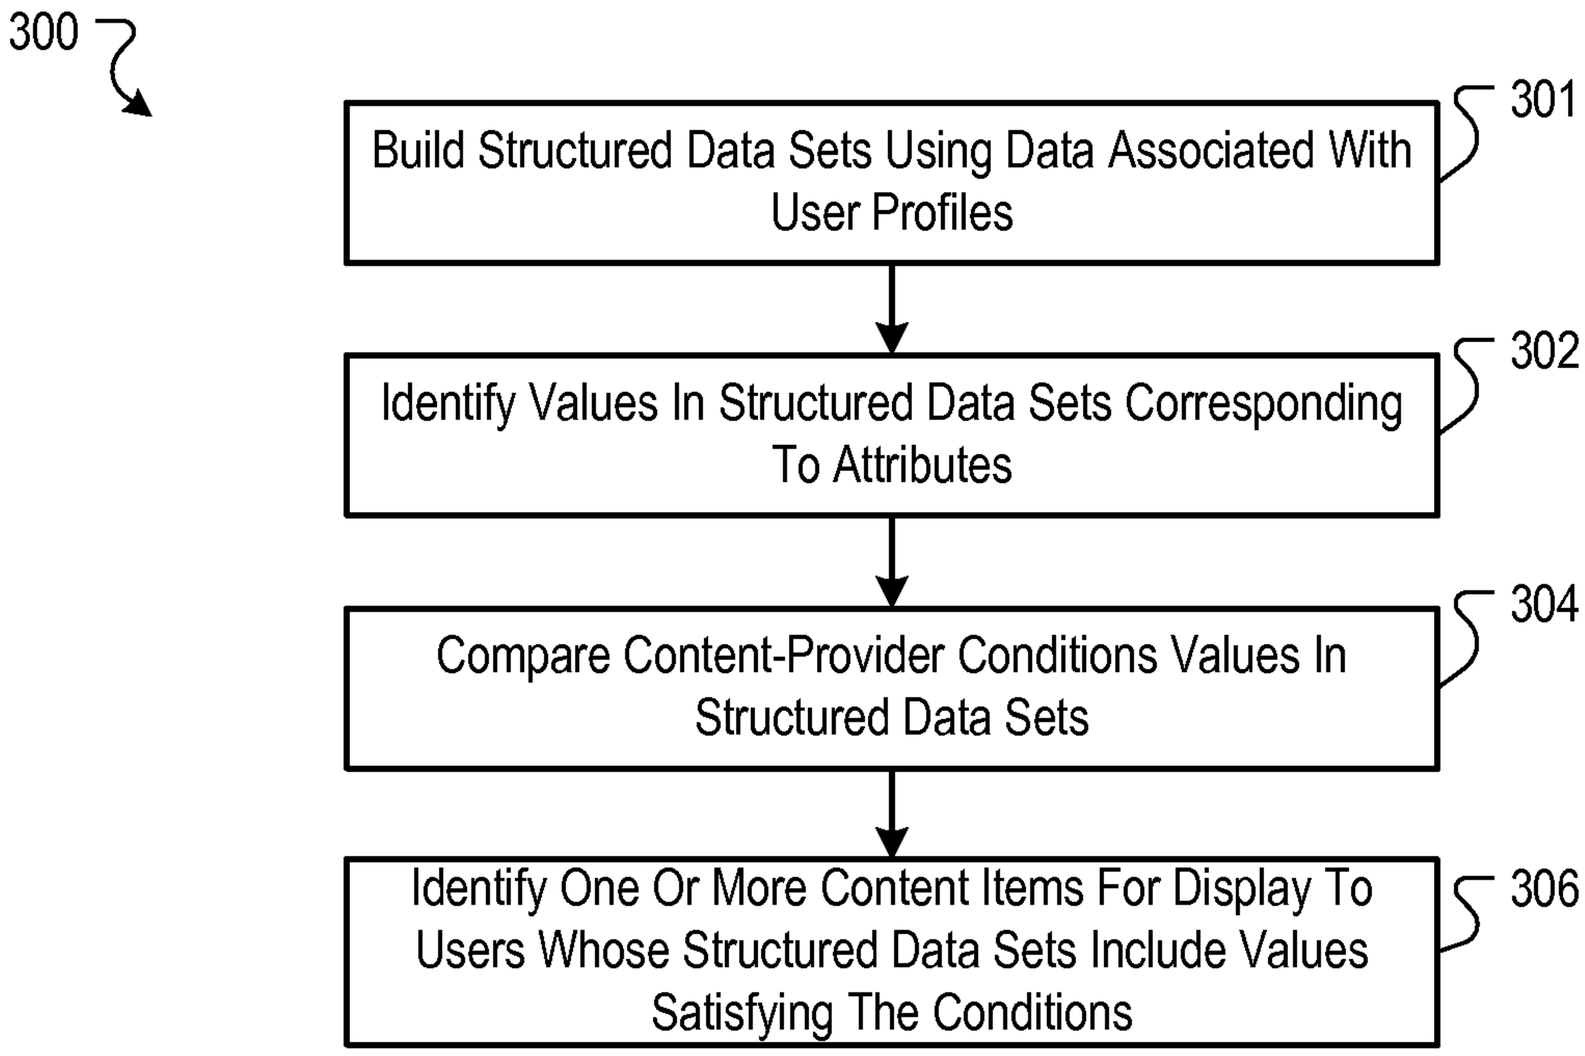

300
Build Structured Data Sets Using Data Associated With User Profiles
301
Identify Values In Structured Data Sets Corresponding To Attributes
302
Compare Content-Provider Conditions Values In Structured Data Sets
304
Identify One Or More Content Items For Display To Users Whose Structured Data Sets Include Values Satisfying The Conditions
306

FIG. 3

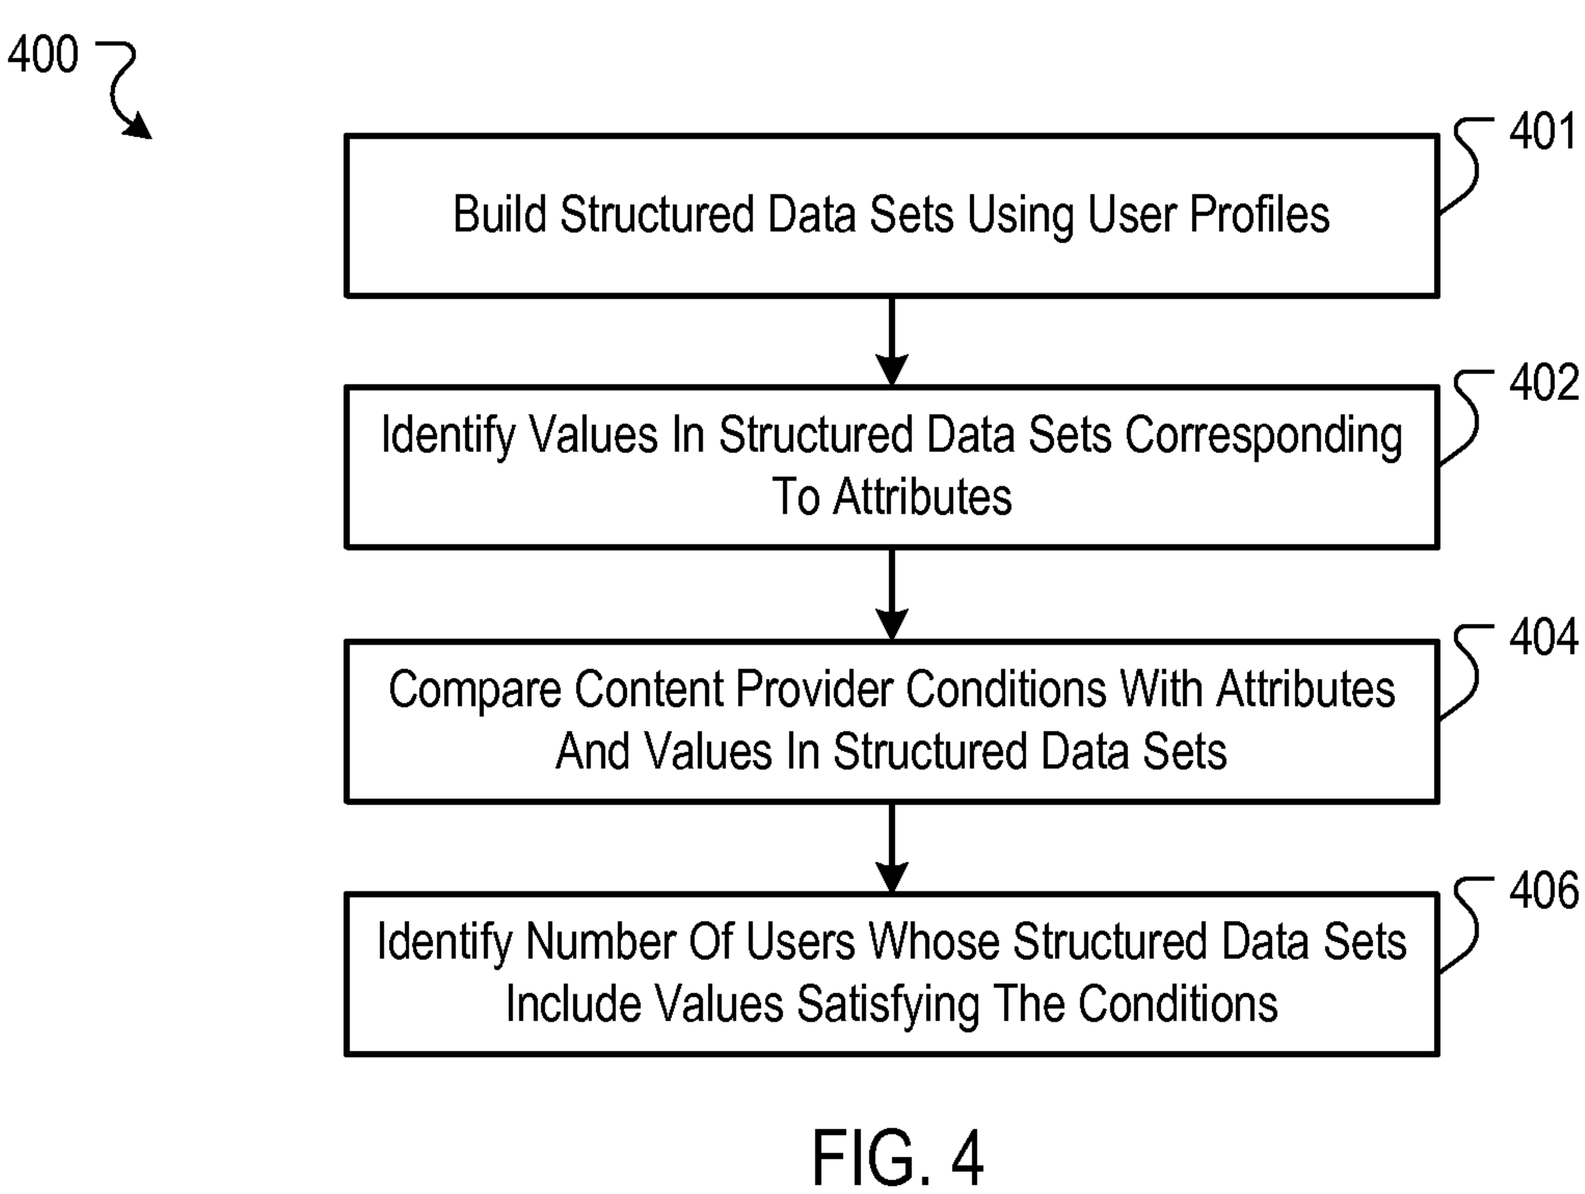

400
Build Structured Data Sets Using User Profiles
401
Identify Values In Structured Data Sets Corresponding To Attributes
402
Compare Content Provider Conditions With Attributes And Values In Structured Data Sets
404
Identify Number Of Users Whose Structured Data Sets Include Values Satisfying The Conditions
406

FIG. 4

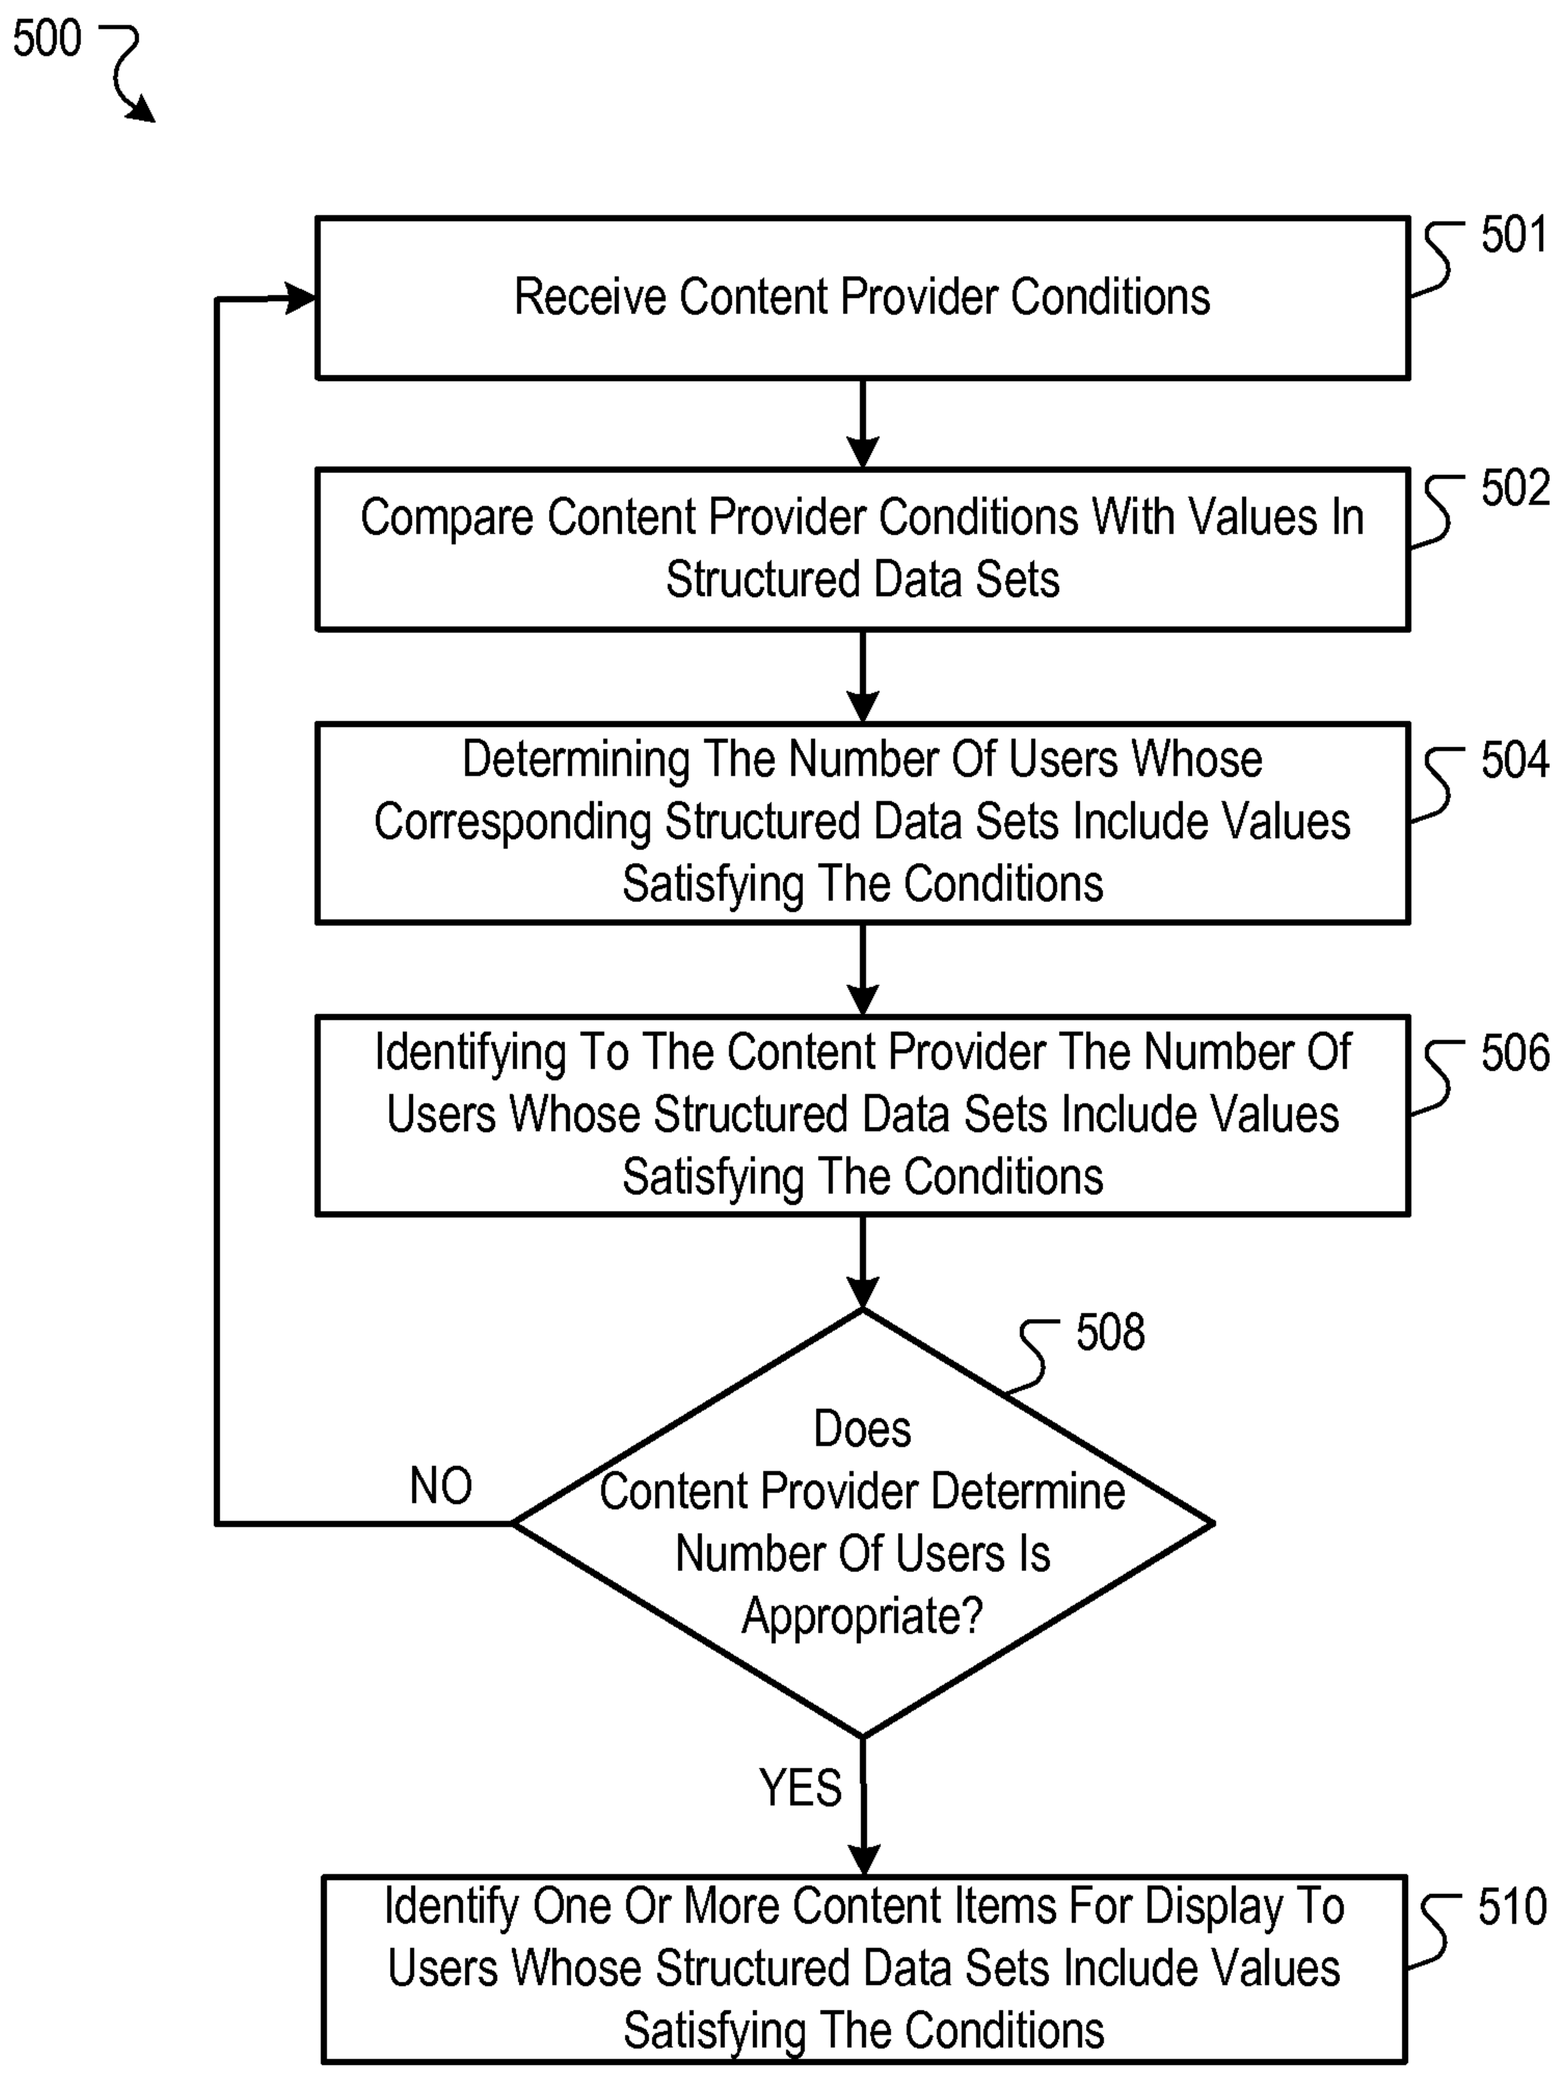

500
Receive Content Provider Conditions
501
Compare Content Provider Conditions With Values In Structured Data Sets
502
Determining The Number Of Users Whose Corresponding Structured Data Sets Include Values Satisfying The Conditions
504
Identifying To The Content Provider The Number Of Users Whose Structured Data Sets Include Values Satisfying The Conditions
506
508
Does Content Provider Determine Number Of Users Is Appropriate?
NO
YES
Identify One Or More Content Items For Display To Users Whose Structured Data Sets Include Values Satisfying The Conditions
510

FIG. 5

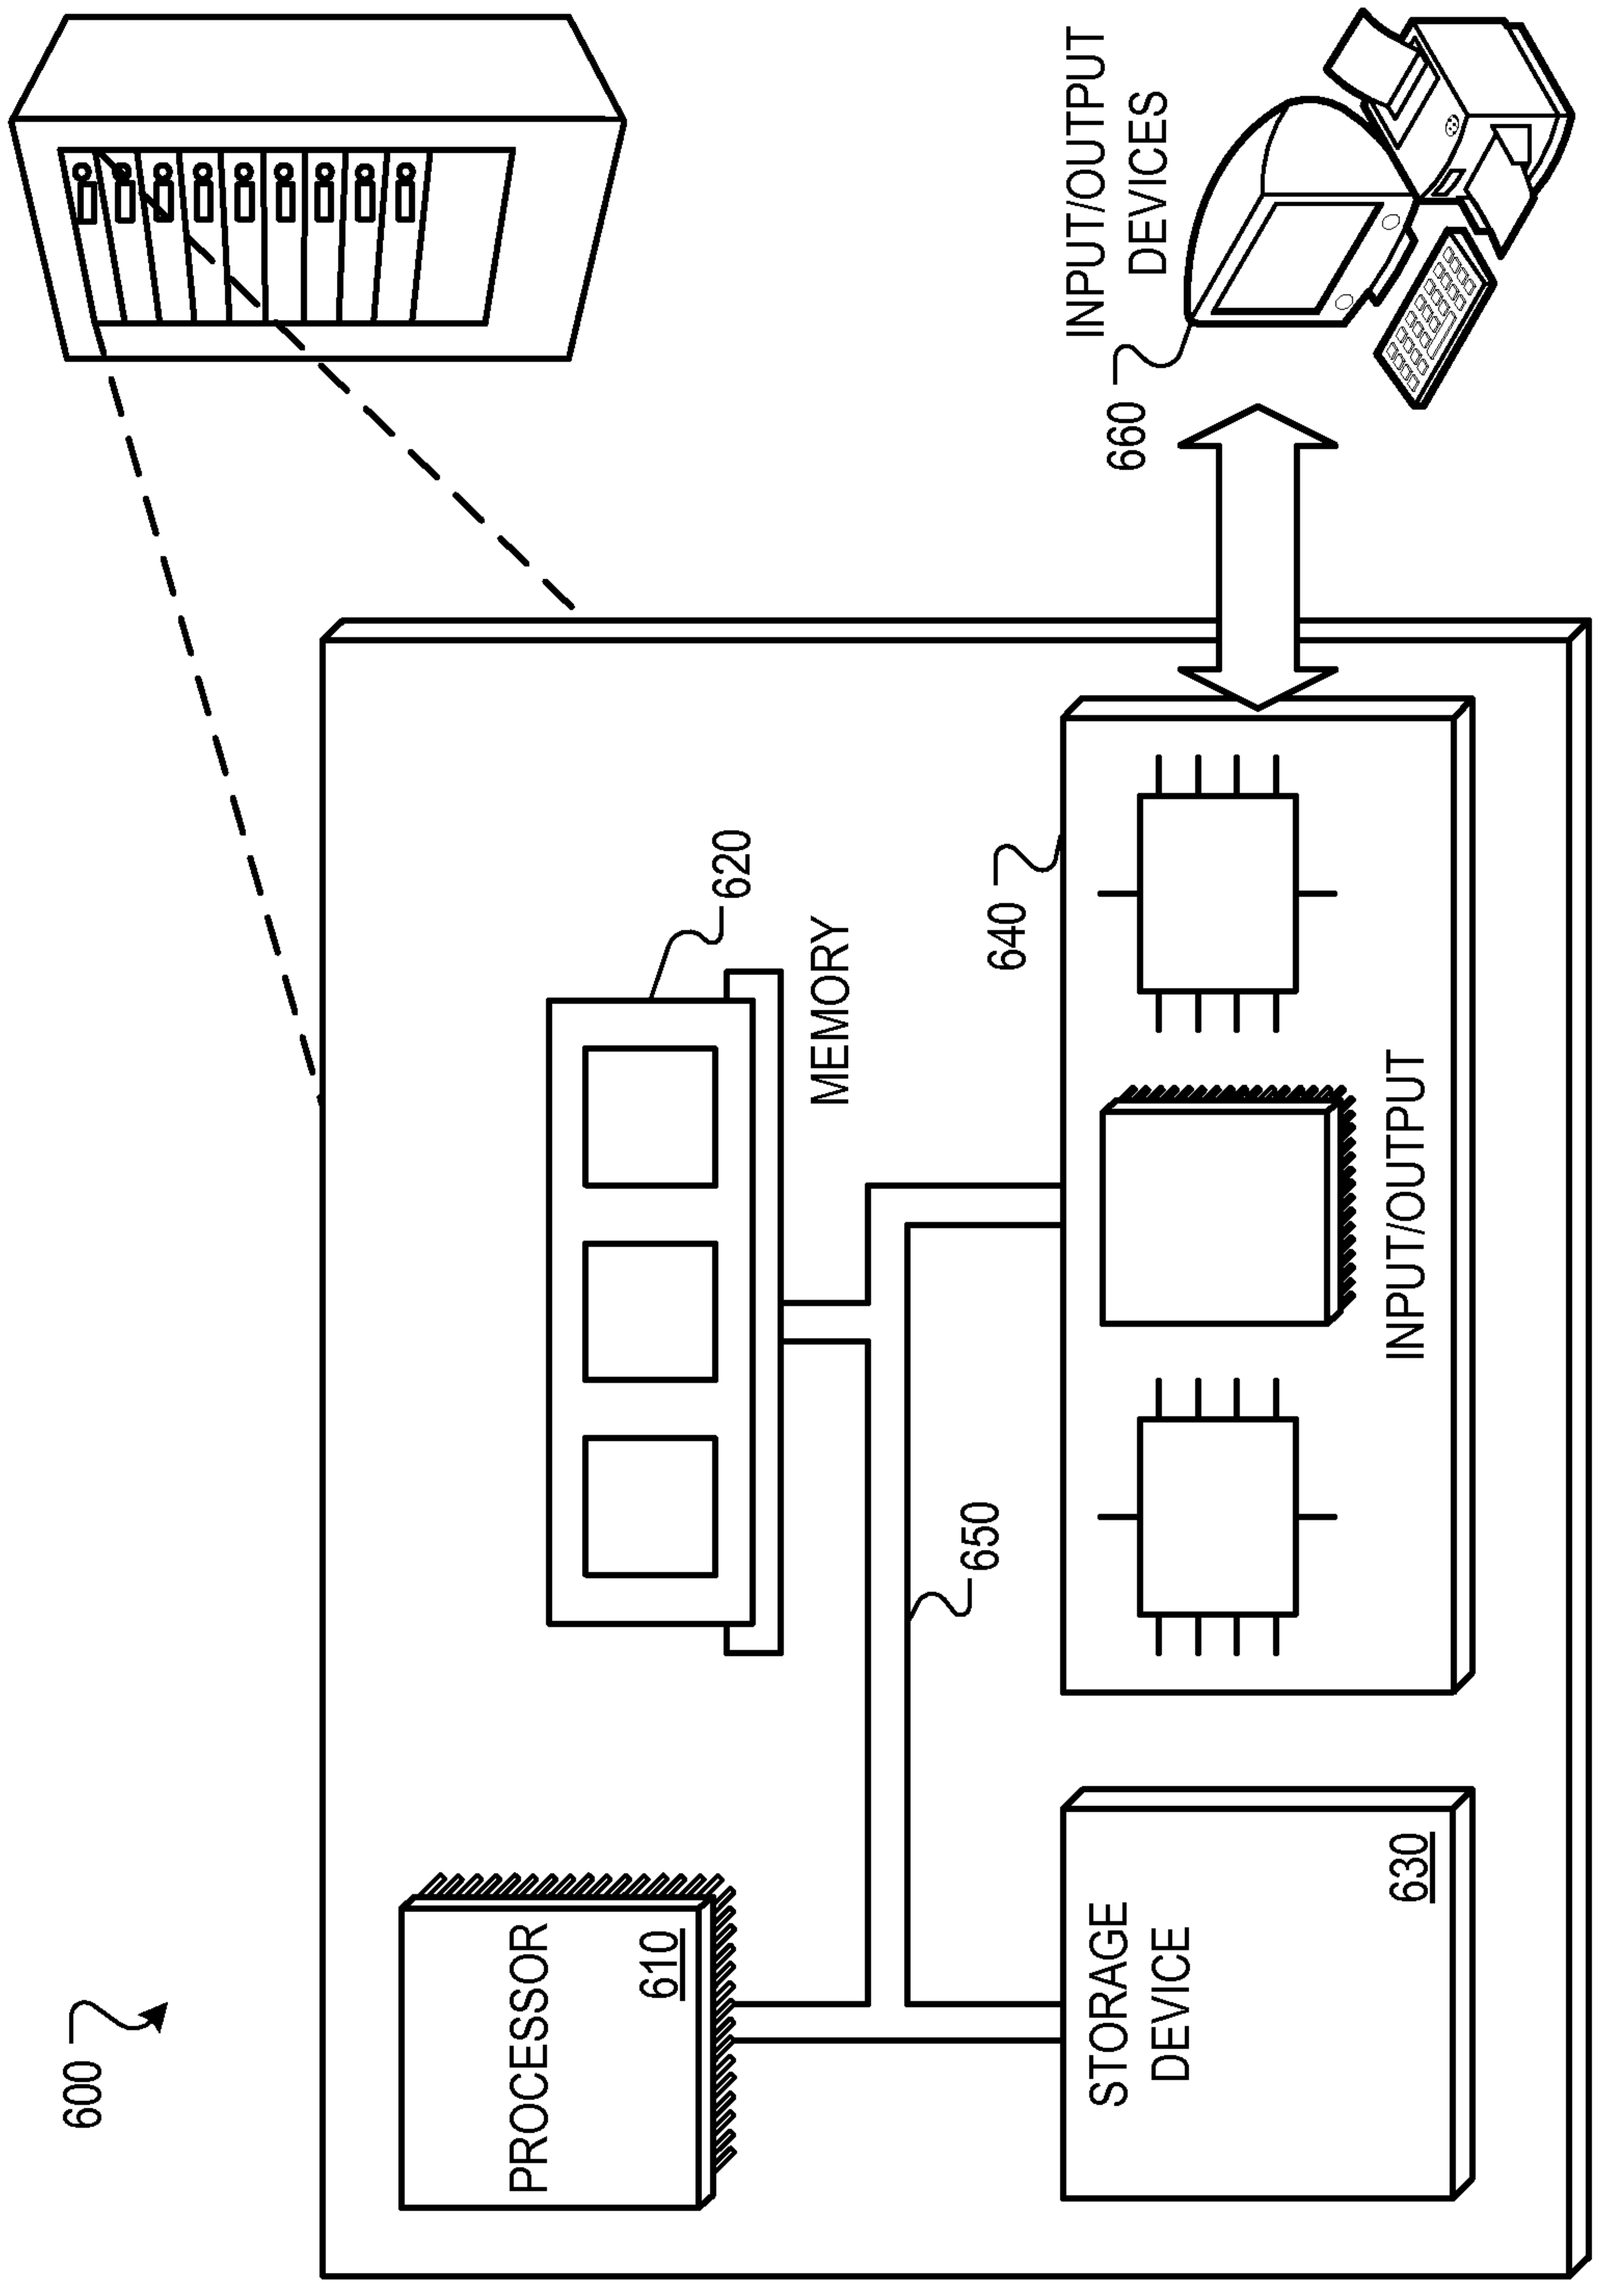
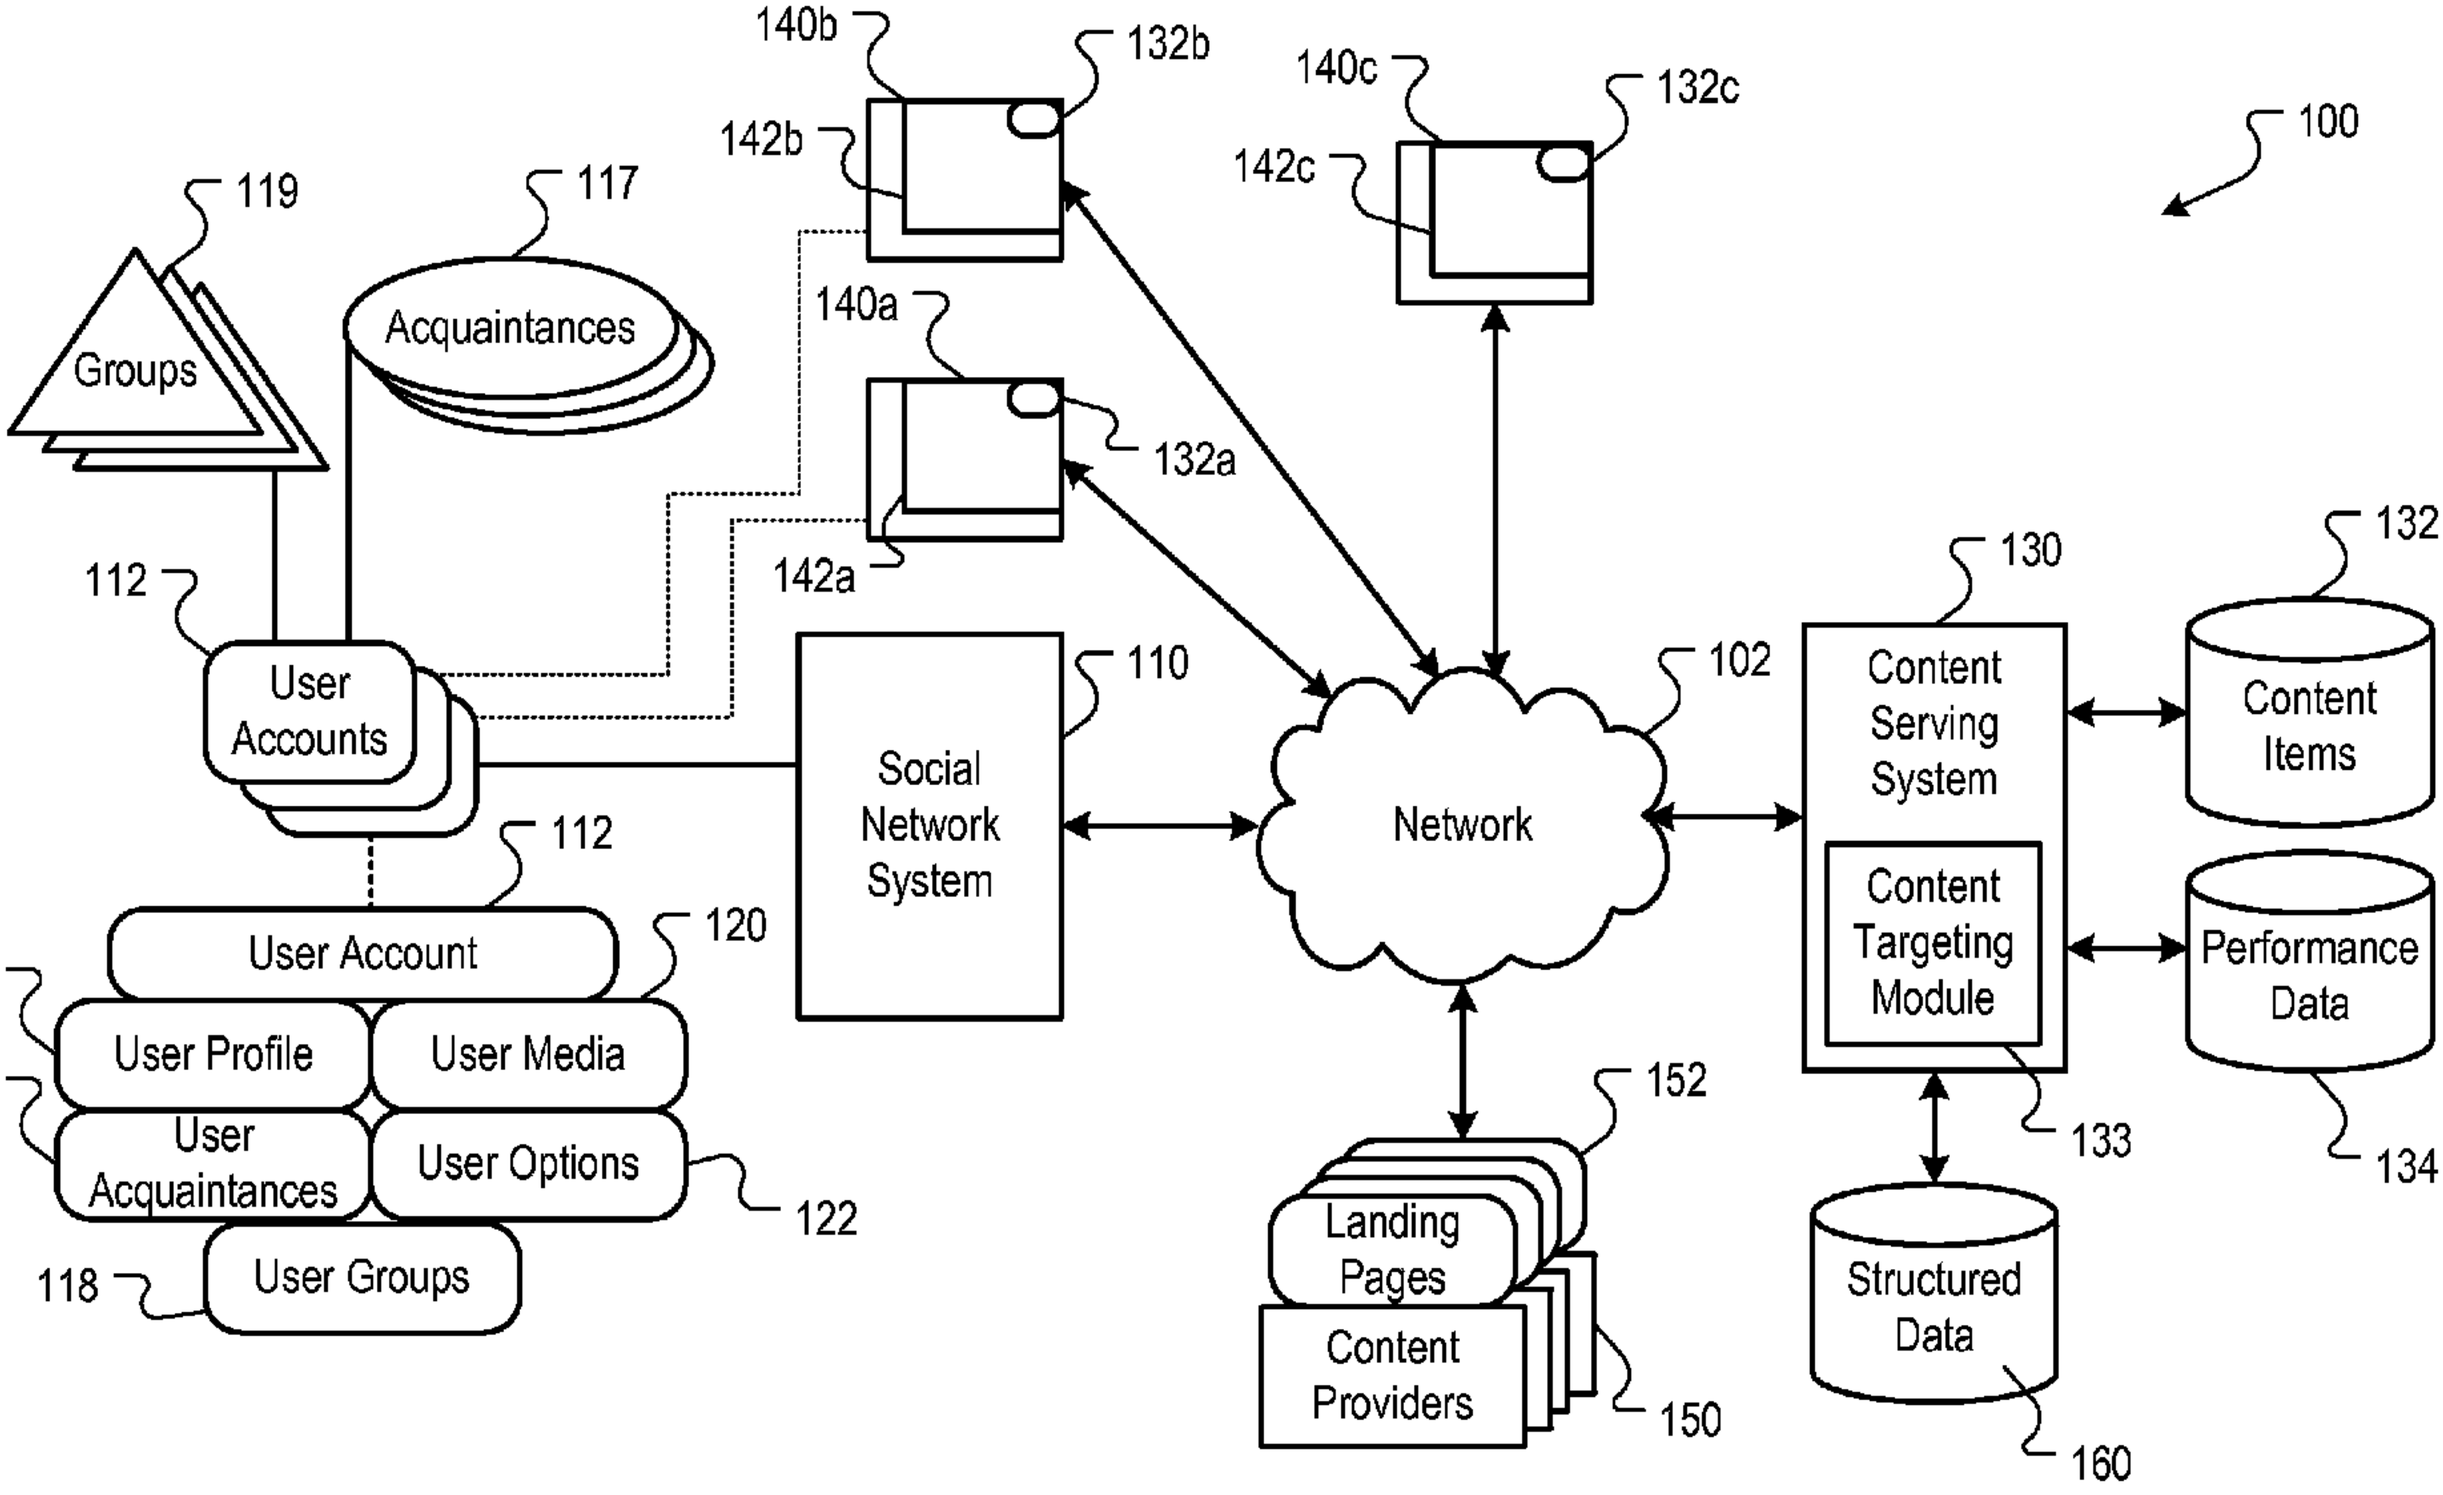

600
PROCESSOR
610
MEMORY
620
650
640
INPUT/OUTPUT
660
INPUT/OUTPUT DEVICES
STORAGE DEVICE
630

FIG. 6


140b
132b
142b
140c
132c
142c
100
119
117
Groups
Acquaintances
140a
132a
142a
112
User Accounts
110
Social Network System
102
Network
130
Content Serving System
132
Content Items
Content Targeting Module
Performance Data
133
134
112
User Account
120
114
User Profile
User Media
116
User Acquaintances
User Options
122
118
User Groups
152
Landing Pages
Content Providers
150
Structured Data
160