

(19) World Intellectual Property Organization  
International Bureau



(43) International Publication Date  
21 July 2005 (21.07.2005)


PCT

(10) International Publication Number  
WO 2005/065229 A2

- (51) International Patent Classification: **Not classified** (US). SAHAMI, Mehran [US/US]; 3228 Hoover Street, Redwood City, CA 94063 (US).
- (21) International Application Number: PCT/US2004/042912 (74) Agent: POKOTYLO, John, C.; Straub & Pokotylo, 620 Tinton Avenue, Bldg. B, 2nd Floor, Tinton Falls, NJ 07724-3260 (US).
- (22) International Filing Date: 21 December 2004 (21.12.2004) (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, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 10/750,363 31 December 2003 (31.12.2003) US (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, 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).
- (71) Applicant (for all designated States except US): **GOOGLE, INC.** [US/US]; 1600 Amphitheatre Parkway, Mountain View, CA 94043 (US).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **BHARAT, Krishna** [IN/US]; 1458 Sierre Creek Way, San Jose, CA 95132 (US). **LAWRENCE, Stephen** [AU/US]; 2400 W. El Camino Real, Apt. 204, Mountain View, CA 94040

[Continued on next page]

(54) Title: GENERATING USER INFORMATION FOR USE IN TARGETED ADVERTISING



(57) **Abstract:** User profile information for a user may be determined by (a) determining initial user profile information for the user, (b) inferring user profile information for the user, and (c) determining the user profile information for the user using both the initial user profile information and the inferred user profile information. Initial user profile information for the user may be determined using past search queries submitted by the user, and/or past document selections by the user. User profile information for the user may be inferred by (a) defining a node for each of a number of documents and the user, (b) adding edges between nodes if there is an association between the nodes to define a graph, and (c) inferring user profile information for the user using a topology of the graph and user profile information of other documents. Similarly, user profile information for a document may be determined by (a) determining initial user profile information for the document, (b) inferring user profile information for the document, and (c) determining the user profile information for the document using both the initial user profile information and the inferred user profile information. The initial user profile information for the document may be determined using content information from the document, and/or document meta information. User profile information for the document may be inferred by (a) defining a node for each of a number of documents and for each of a number of users, (b) adding edges between nodes if there is an association between the nodes to define a graph, and (c) inferring user profile information for the document using a topology of the graph and user profile information of users and of other documents. Document, user, and/or ad user profile information may be used when serving ads.



WO 2005/065229 A2

**Declarations under Rule 4.17:**

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations 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, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, 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, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for the following designations 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, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, 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, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LT, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (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.

**GENERATING USER INFORMATION FOR USE IN TARGETED ADVERTISING****§ 1. BACKGROUND OF THE INVENTION**5           **§ 1.1 FIELD OF THE INVENTION**

The present invention concerns advertising. In particular, the present invention concerns determining particularly relevant advertisements or advertisement creatives to serve for a user request, such as a search query or document request for example.

10

**§ 1.2 BACKGROUND INFORMATION**

Advertising using traditional media, such as television, radio, newspapers and magazines, is well known. Unfortunately, even when armed with demographic studies and  
15 entirely reasonable assumptions about the typical audience of various media outlets, advertisers recognize that much 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  
20 and services offered over the Internet as a potentially powerful way to advertise.

Advertisers have developed several strategies in an attempt to maximize the value of such advertising. In one strategy, advertisers use popular presences or means for providing interactive media or services (referred to as “Websites” in the specification without loss of generality) as conduits to reach a large audience. Using this first approach, an advertiser may  
25 place ads on the home page of the New York Times Website, or the USA Today Website, for example. In another strategy, an advertiser may attempt to target its ads to narrower niche audiences, thereby increasing the likelihood of a positive response by the audience. For example, an agency promoting tourism in the Costa Rican rainforest might place ads on the ecotourism-travel subdirectory of the Yahoo Website. An advertiser will normally determine  
30 such targeting manually.

Regardless of the strategy, Website-based ads (also referred to as “Web ads”) are typically 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 (referred to as a “viewer” or “user” in the specification 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 Website. 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 rendered) is commonly referred to as the "click-through rate" of the ad.

A "conversion" is said to occur when a user consummates a transaction related to a previously served ad. What constitutes a conversion may vary from case to case and can be determined in a variety of ways. For example, it may be the case that a conversion occurs when a user clicks on an ad, is referred to the advertiser's web page, and consummates a purchase there before leaving that Web page. Alternatively, a conversion may be defined as a user being shown an ad, and making a purchase on the advertiser's Web page within a predetermined time (e.g., seven days). In yet another alternative, a conversion may be defined by an advertiser to be any measurable/observable user action such as, for example, downloading a white paper, navigating to at least a given depth of a Website, viewing at least a certain number of Web pages, spending at least a predetermined amount of time on a Website or Web page, registering on a Website, etc. Often, if user actions don't indicate a consummated purchase, they may indicate a sales lead, although user actions constituting a conversion are not limited to this. Indeed, many other definitions of what constitutes a conversion are possible. The ratio of the number of conversions to the number of impressions of the ad (i.e., the number of times an ad is displayed) is commonly referred to as the conversion rate. If a conversion is defined to be able to occur within a predetermined time since the serving of an ad, one possible definition of the conversion rate might only consider ads that have been served more than the predetermined time in the past.

Despite the initial promise of Website-based advertisement, there remain several problems with existing approaches. Although advertisers are able to reach a large audience, they are frequently dissatisfied with the return on their advertisement investment.

Similarly, the hosts of Websites on which the ads are presented (referred to as "Website hosts" or "ad consumers") have the challenge of maximizing ad revenue without impairing their users' experience. Some Website hosts have chosen to place advertising revenues over the interests of users. One such Website is "Overture.com", which hosts a so-called "search engine" service returning advertisements masquerading as "search results" in response to user queries. The Overture.com Website permits advertisers to pay to position an ad for their Website (or a target Website) higher up on the list of purported search results. If such schemes where the advertiser only pays if a user clicks on the ad (i.e., cost-per-click) are implemented, the

advertiser lacks incentive to target their ads effectively, since a poorly targeted ad will not be clicked and therefore will not require payment. Consequently, high cost-per-click ads show up near or at the top, but do not necessarily translate into real revenue for the ad publisher because viewers don't click on them. Furthermore, ads that viewers would click on are further down the list, or not on the list at all, and so relevancy of ads is compromised.

Search engines, such as Google for example, have enabled advertisers to target their ads so that they will be rendered with a search results page and so that they will be relevant, presumably, to the query that prompted the search results page. Other targeted advertising systems, such as those that target ads based on e-mail information (See, e.g., the systems described in U.S. Patent Application Serial No. 10/452,830 (incorporated herein by reference), titled "SERVING ADVERTISEMENTS USING INFORMATION ASSOCIATED WITH E-MAIL", filed on June 2, 2003 and listing Jeffrey A. Dean, Georges R. Harik and Paul Bucheit as inventors.); or those that target ads based on content (See, e.g., U.S. Patent Application Serial No. 10/375,900 (incorporated herein by reference), titled "SERVING ADVERTISEMENTS BASED ON CONTENT", filed on February 26, 2003 and listing Darrell Anderson, Paul Bucheit, Alex Carobus, Claire Cui, Jeffrey A. Dean, Georges R. Harik, Deepak Jindal, and Narayanan Shivakumar as inventors.) may have similar challenges. That is, advertising systems would like to present advertisements that are relevant to the user requested information in general, and related to the current user interest in particular.

As can be appreciated from the foregoing, targeted advertising systems, such as keyword-targeted advertising or content-targeted advertising provide very useful forms of advertising. However, even online advertising systems that can serve generally relevant advertisements often cannot select relevant advertisements best suited for a particular user.

Accordingly, there is a need to improve the performance of online advertising. More specifically, there is a need to increase the relevancy of ads served for some user request, such as a search query or a document request for example, to the user that submitted the request.

## § 2. SUMMARY OF THE INVENTION

The present invention describes methods and apparatus for determining a match used for scoring an ad. Such methods and apparatus may (a) determine a first match value using (i) user profile information of an ad landing page of the ad and/or user profile information used for targeting the ad, and (ii) user profile information of a user to which the ad will be rendered, (b) determining a second match value using (i) user profile information of an ad landing page of the

ad and/or user profile information used for targeting the ad, and (ii) user profile information of a document with which the ad will be served, and (c) determining the match used for scoring the ad using the first match value and the second match value.

5 The present invention also teaches that in methods and apparatus such as the foregoing, or in other methods and apparatus for targeting ads using user profile information, at least some of the user profile information of the ad landing page of the ad, at least some of the user profile information used for targeting of the ad, at least some of the user profile information of the user, and/or at least some of the user profile information of the document may be inferred.

10 In one embodiment of the present invention, user profile information for a user may be determined by (a) determining initial user profile information for the user, (b) inferring user profile information for the user, and (c) determining the user profile information for the user using both the initial user profile information and the inferred user profile information. Initial user profile information for the user may be determined using past search queries submitted by the user, and/or past document selections by the user.

15 In one embodiment of the present invention, user profile information for the user may be inferred by (a) defining a node for each of a number of documents and the user, (b) adding edges between nodes if there is an association between the nodes to define a graph, and (c) inferring user profile information for the user using a topology of the graph and user profile information of other documents.

20 In one embodiment of the present invention, user profile information for a document may be determined by (a) determining initial user profile information for the document, (b) inferring user profile information for the document, and (c) determining the user profile information for the document using both the initial user profile information and the inferred user profile information. The initial user profile information for the document may be determined  
25 using content information from the document, and/or document meta information.

In one embodiment of the present invention, user profile information for the document may be inferred by (a) defining a node for each of a number of documents and for each of a number of users, (b) adding edges between nodes if there is an association between the nodes to define a graph, and (c) inferring user profile information for the document using a topology of  
30 the graph and user profile information of users and of other documents.

### § 3. BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a high-level diagram showing parties or entities that can interact with an advertising system.

5 Figure 2 is a bubble chart of an exemplary advertising environment in which, or with which, the present invention may operate.

Figure 3 is a bubble chart of operations, consistent with the present invention, which may be used in, or with, an online advertising environment, such as the one in Figure 2.

10 Figure 4 illustrates an exemplary data structure for storing user profile information that may be generated, updated and/or used in a manner consistent with the present invention.

Figure 5 illustrates possible components of an overall match value which may be a factor in an ad score.

Figure 6 is a flow diagram of an exemplary method 600 that may be used to determine document UPI in a manner consistent with the present invention.

15 Figure 7 is a flow diagram of an exemplary method 700 that may be used to determine user UPI in a manner consistent with the present invention.

Figure 8 is a flow diagram of an exemplary method 800 that may be used to determine an initial or baseline document UPI in a manner consistent with the present invention.

20 Figure 9 is a flow diagram of an exemplary method 900 that may be used to determine an initial or baseline user UPI in a manner consistent with the present invention.

Figure 10 illustrates how users and documents can be associated.

Figure 11 is a flow diagram of an exemplary method 1100 that may be used to associated users and/or documents in a manner consistent with the present invention.

25 Figure 12 is a block diagram of apparatus that may be used to effect at least some of the various operations that may be performed, and to store information that may be used and/or generated, in a manner consistent with the present invention.

### § 4. DETAILED DESCRIPTION

30 The present invention may involve novel methods, apparatus, message formats and/or data structures for determining user profile information and using such determined user profile information for ad serving. The following description is presented to enable one skilled in the art to make and use the invention, and is provided in the context of particular applications and their requirements. Various modifications to the disclosed embodiments will be apparent to

those skilled in the art, and the general principles set forth below may be applied to other embodiments and applications. Thus, the present invention is not intended to be limited to the embodiments shown and the inventors regard their invention as any patentable subject matter described.

5           One possible method to improve ad targeting is for ad targeting systems to obtain and use user profiles. For example, user profiles may be determined using information voluntarily given by users (e.g., when they subscribe to a service). This user attribute information may then be matched against advertiser specified attributes of the ad (e.g., targeting criteria). Unfortunately, user profile information is not always available since many Websites (e.g., search engines) do not require subscription or user registration. Moreover, even when available, the user profile may be incomplete (e.g., because the information given at the time of subscription may be limited to what is needed for the service and hence not comprehensive, because of privacy considerations, etc.). Furthermore, advertisers may need to manually define user profile targeting information. In addition, even if user profile information is available, advertisers may not be able to use this information to target ads effectively.

10           In the following, environments in which, or with which, the present invention may operate are described in § 4.1. Then, exemplary embodiments of the present invention are described in § 4.2. Some alternatives and refinements are described in § 4.3. Finally, some conclusions regarding the present invention are set forth in § 4.4.

## 20           **§ 4.1 ENVIRONMENTS IN WHICH, OR WITH WHICH, THE PRESENT INVENTION MAY OPERATE**

### 25           **§ 4.1.1 EXEMPLARY ADVERTISING ENVIRONMENT**

Figure 1 is a high level diagram of an advertising environment. The environment may include an ad entry, maintenance and delivery system 120. Advertisers 110 may directly, or indirectly, enter, maintain, and track ad information in the system 120. The ads may be in the form of graphical ads such as so-called banner ads, text only ads, image ads, audio ads, video ads, ads combining one of more of any of such components, etc. The ads may also include embedded information, such as a link, and/or machine executable instructions. Ad consumers 130 may submit requests for ads to, accept ads responsive to their request from, and provide usage information to, the system 120. Although not shown, other entities may provide usage information (e.g., whether or not a conversion or click-through related to the ad occurred) to the



system 120. This usage information may include measured or observed user behavior related to ads that have been served.

One example of an ad consumer 130 is a general content server that receives requests for content (e.g., articles, discussion threads, music, video, graphics, search results, web page listings, etc.), and retrieves the requested content in response to, or otherwise services, the request. The content server may submit a request for ads to the system 120. Such an ad request may include a number of ads desired. The ad request may also include content request information. This information may include the content itself (e.g., page), a category corresponding to the content or the content request (e.g., arts, business, computers, arts-movies, arts-music, etc.), part or all of the content request, content age, content type (e.g., text, graphics, video, audio, mixed media, etc.), geolocation information, etc.

The content server may combine the requested content with one or more of the advertisements provided by the system 120. This combined information including the content and advertisement(s) is then forwarded towards the end user that requested the content, for presentation to the user. Finally, the content server may transmit information about the ads and how, when, and/or where the ads are to be rendered (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the system 120. Alternatively, or in addition, such information may be provided back to the system 120 by some other means.

Another example of an ad consumer 130 is a search engine. A search engine may receive queries for search results. In response, the search engine may retrieve relevant search results (e.g., from an index of Web pages). An exemplary search engine is described in the article S. Brin and L. Page, "The Anatomy of a Large-Scale Hypertextual Search Engine," Seventh International World Wide Web Conference, Brisbane, Australia and in U.S. Patent No. 6,285,999 (both incorporated herein by reference). Such search results may include, for example, lists of Web page titles, snippets of text extracted from those Web pages, and hypertext links to those Web pages, and may be grouped into a predetermined number of (e.g., ten) search results.

The search engine may submit a request for ads to the system 120. The request may include a number of ads desired. This number may depend on the search results, the amount of screen or page space occupied by the search results, the size and shape of the ads, etc. In one embodiment, the number of desired ads will be from one to ten, and preferably from three to five. The request for ads may also include the query (as entered or parsed), information based on the query (such as geolocation information, whether the query came from an affiliate and an

identifier of such an affiliate), and/or information associated with, or based on, the search results. Such information may include, for example, identifiers related to the search results (e.g., document identifiers or “docIDs”), scores related to the search results (e.g., information retrieval (“IR”) scores such as dot products of feature vectors corresponding to a query and a document, Page Rank scores, and/or combinations of IR scores and Page Rank scores), snippets of text extracted from identified documents (e.g., Web Pages), full text of identified documents, feature vectors of identified documents, etc.

The search engine may combine the search results with one or more of the advertisements provided by the system 120. This combined information including the search results and advertisement(s) is then forwarded towards the user that requested the content, for presentation to the user. Preferably, the search results are maintained as distinct from the ads, so as not to confuse the user between paid advertisements and presumably neutral search results.

Finally, the search engine may transmit information about the ad and when, where, and/or how the ad was to be rendered (e.g., position, click-through or not, impression time, impression date, size, conversion or not, etc.) back to the system 120. Alternatively, or in addition, such information may be provided back to the system 120 by some other means.

#### **§ 4.1.2 EXEMPLARY AD ENTRY, MAINTENANCE AND DELIVERY ENVIRONMENT**

Figure 2 illustrates an exemplary ad system 120' in which, or with which, the present invention may be used. The exemplary ad system 120' may include an inventory system 210 and may store ad information 205 and usage information 245. The exemplary system 120' may support ad information entry and management operations 215, campaign (e.g., targeting) assistance operations 220, accounting and billing operations 225, ad serving operations 230, relevancy determination operations 235, optimization operations 240, relative presentation attribute assignment (e.g., position ordering) operations 250, fraud detection operations 255, and result interface operations 260.

Advertisers 110 may interface with the system 120' via the ad information entry and management operations 215 as indicated by interface 216. Ad consumers 130 may interface with the system 120' via the ad serving operations 230 as indicated by interface 231. Ad consumers 130 and/or other entities (not shown) may also interface with the system 120' via results interface operations 260 as indicated by interface 261.

An advertising program may include information concerning accounts, campaigns, creatives, targeting, etc. The term “account” relates to information for a given advertiser (e.g., a

unique email address, a password, billing information, etc.). A “campaign” or “ad campaign” refers to one or more groups of one or more advertisements, and may include a start date, an end date, budget information, geo-targeting information, syndication information, etc. For example, Honda may have one advertising campaign for its automotive line, and a separate advertising campaign for its motorcycle line. The campaign for its automotive line may have one or more ad groups, each containing one or more ads. Each ad group may include a set of keywords, and a maximum cost bid (cost per click-through, cost per conversion, etc.). Alternatively, or in addition, each ad group may include an average cost bid (e.g., average cost per click-through, average cost per conversion, etc.). Therefore, a single maximum cost bid and/or a single average cost bid may be associated with one or more keywords. As stated, each ad group may have one or more ads or “creatives” (That is, ad content that is ultimately rendered to an end user.). Naturally, the ad information 205 may include more or less information, and may be organized in a number of different ways.

The ad information 205 can be entered and managed via the ad information entry and management operations 215. Campaign (e.g., targeting) assistance operations 220 can be employed to help advertisers 110 generate effective ad campaigns. For example, the campaign assistance operations 220 can use information provided by the inventory system 210, which, in the context of advertising for use with a search engine, may track all possible ad impressions, ad impressions already reserved, and ad impressions available for given keywords.

The ad serving operations 230 may service requests for ads from ad consumers 130. The ad serving operations 230 may use relevancy determination operations 235 to determine candidate ads for a given request. The ad serving operations 230 may then use optimization operations 240 to select a final set of one or more of the candidate ads. The ad serving operations 230 may then use relative presentation attribute assignment operations 250 to order the presentation of the ads to be returned. The accounting/billing operations 225 may be used to track charges related to the serving of advertisements and to bill advertisers. The fraud detection operations 255 can be used to reduce fraudulent use of the advertising system (e.g., by advertisers), such as through the use of stolen credit cards. Finally, the results interface operations 260 may be used to accept result information (from the ad consumers 130 or some other entity) about an ad actually served, such as whether or not click-through occurred, whether or not conversion occurred (e.g., whether the sale of an advertised item or service was initiated or consummated within a predetermined time from the rendering of the ad), etc. Such results information may be accepted at interface 261 and may include information to identify the ad and time the ad was served, as well as the associated result.

Various aspects of the present invention may be used with relevancy determination operations 235, relative presentation attribute ordering operations 250, and/or ad serving operations 230.

### 5 § 4.1.3 EXEMPLARY NETWORK ENVIRONMENT

Figure 3 illustrates an exemplary network environment 300 in which the present invention may be used. The exemplary network environment 300 may include one or more client devices 310, each having a browser 320 (or some other information requesting and rendering means). The client devices 310 can request documents 335 served by one or more content servers 320 and can search content included in those or other documents using one or more search engines 340. An ad server 350 can serve one or more ads 355. The one or more ads 355 served may be relevant to documents served by the content server 330 (and/or the request for such documents) and/or relevant to search results generated by search engine 340 (and/or the search query). User information (e.g., about an individual user or group of users associated with one or more client devices 310) 365 may be stored at one or more information servers 360. Alternatively, or in addition, user information (not shown) may be stored at one or more content servers 330, one or more search engines 340, and/or one or more ad servers 350. The one or more client devices 310, content servers 330, search engines 340, ad servers 350, and/or user information servers 360 may exchange information with one another via one or more networks 370. The one or more networks 370 may be the Internet and the servers and search engines may be computers. The user information (e.g., user profile database), an information index, and an advertisement index need not be separate repositories -- they may be stored in a single repository. Moreover, these forms of data may be interleaved (for example, the advertisement database may include "user feature tags" on the ads so that only some subsets are retrieved for particular types of users).

### § 4.1.4 DEFINITIONS

30 Online ads, such as those used in the exemplary systems described above with reference to Figures 1, 2, 3 or any other system, may have various intrinsic features. Such features may be specified by an application and/or an advertiser. These features are referred to as "ad features" below. For example, in the case of a text ad, ad features may include a title line, ad text, and an embedded link. In the case of an image ad, ad features may include images, executable code,

and an embedded link. Depending on the type of online ad, ad features may include one or more of the following: text, a link, an audio file, a video file, an image file, executable code, embedded information, etc.

When an online ad is served, one or more parameters may be used to describe how, when, and/or where the ad was served. These parameters are referred to as “serving parameters” below. Serving parameters may include, for example, one or more of the following: features of (including information on) a page on which the ad was served, a search query or search results associated with the serving of the ad, a user characteristic (e.g., their geolocation, the language used by the user, the type of browser used, previous page views, previous behavior), a host or affiliate site (e.g., America Online, Google, Yahoo) that initiated the request, an absolute position of the ad on the page on which it was served, a position (spatial or temporal) of the ad relative to other ads served, an absolute size of the ad, a size of the ad relative to other ads, a color of the ad, a number of other ads served, types of other ads served, time of day served, time of week served, time of year served, etc. Naturally, there are other serving parameters that may be used in the context of the invention.

Although serving parameters may be extrinsic to ad features, they may be associated with an ad as serving conditions or constraints. When used as serving conditions or constraints, such serving parameters are referred to simply as “serving constraints” (or “targeting criteria”). For example, in some systems, an advertiser may be able to target the serving of its ad by specifying that it is only to be served on weekdays, no lower than a certain position, only to users in a certain geolocation, etc. As another example, in some systems, an advertiser may specify that its ad is to be served only if a page or search query includes certain keywords or phrases. As yet another example, in some systems, an advertiser may specify that its ad is to be served only if a document being served includes certain topics or concepts, or falls under a particular cluster or clusters, or some other classification or classifications. “Geolocation information” may include information specifying one or more of one or more countries, one or more (inter-country) regions, one or more states, one or more metro areas, one or more cities, one or more towns, one or more boroughs, one or more areas with common zip codes, one or more areas with common telephone area codes, one or more areas served by common cable head end stations, one or more areas served by common network access points or nodes, etc. It may include latitude and/or longitude, or a range thereof. It may include information, such as an IP address, from which a user location can be estimated.

“Ad information” may include any combination of ad features, ad serving constraints, information derivable from ad features or ad serving constraints (referred to as “ad derived

information”), and/or information related to the ad (referred to as “ad related information”), as well as an extension of such information (e.g., information derived from ad related information).

A “document” is to be broadly interpreted to include any machine-readable and machine-storable work product. A document may be a file, a combination of files, one or more files with embedded links to other files, etc. The files may be of any type, such as text, audio, image, video, etc. Parts of a document to be rendered to an end user can be thought of as “content” of the document. A document may include “structured data” containing both content (words, pictures, etc.) and some indication of the meaning of that content (for example, e-mail fields and associated data, HTML tags and associated data, etc.) Ad spots in the document may be defined by embedded information or instructions. In the context of the Internet, a common document is a Web page. Web pages often include content and may include embedded information (such as meta information, hyperlinks, etc.) and/or embedded instructions (such as Javascript, etc.). In many cases, a document has a unique, addressable, storage location and can therefore be uniquely identified by this addressable location. A universal resource locator (URL) is a unique address used to access information on the Internet.

“Document information” may include any information included in the document, information derivable from information included in the document (referred to as “document derived information”), and/or information related to the document (referred to as “document related information”), as well as an extensions of such information (e.g., information derived from related information). An example of document derived information is a classification based on textual content of a document. Examples of document related information include document information from other documents with links to the instant document, as well as document information from other documents to which the instant document links.

“E-mail information” may include any information included in an e-mail (also referred to as “internal e-mail information”), information derivable from information included in the e-mail and/or information related to the e-mail, as well as extensions of such information (e.g., information derived from related information). An example of information derived from e-mail information is information extracted or otherwise derived from search results returned in response to a search query composed of terms extracted from an e-mail subject line. Examples of information related to e-mail information include e-mail information about one or more other e-mails sent by the same sender of a given e-mail, or user information about an e-mail recipient. Information derived from or related to e-mail information may be referred to as “external e-mail information.”

Content from a document may be rendered on a “content rendering application or device”. Examples of content rendering applications include an Internet browser (e.g., Explorer or Netscape), a media player (e.g., an MP3 player, a Realnetworks streaming audio file player, etc.), a viewer (e.g., an Adobe Acrobat pdf reader), etc.

5 A “content owner” is a person or entity that has some property right in the content of a document. A content owner may be an author of the content. In addition, or alternatively, a content owner may have rights to reproduce the content, rights to prepare derivative works of the content, rights to display or perform the content publicly, and/or other proscribed rights in the content. Although a content server might be a content owner in the content of the documents it  
10 serves, this is not necessary.

“User profile information” (also referred to as “UPI”) may include any information about an individual user or a group of users. Such information may be provided by the user, provided by a third party authorized to release user information, and/or derived from user actions. Certain user information can be deduced or presumed using other user information of the same user  
15 and/or user information of other users. UPI may be associated with various entities. “User UPI” is user profile information associated with a particular user or a group of users. “Document UPI” is user profile information associated with a document. For example, a document UPI may be a composite representation of various user UPI’s of users that have requested or visited the document. “Ad landing page UPI” is a particular type of document UPI and may be a composite  
20 representation of various users that have visited, or consummated a purchase on, a particular Web page, perhaps in response to their selecting (e.g., clicking) an ad. “Ad targeting UPI” may include user profile serving constraints. For example, an ad for prostate cancer screening might be limited to user profiles having the attribute “male” and “age 45 and over.”

Various exemplary embodiments of the present invention are now described in § 4.2.  
25

## § 4.2 EXEMPLARY EMBODIMENTS

### § 4.2.1 EXEMPLARY USER INFORMATION

30 Figure 4 is an exemplary data structure 400 that may be used to store user profile information (UPI) in a manner consistent with the present invention. Although not shown, that data structure 400 may include an identifier. Such an identifier may be used to associate the UPI with a particular user, a group of users, a document, an ad, an ad landing page, etc. The UPI may include one or more collections of information 410 related to one or more UPI attributes  
35 420. Each attribute 420 may have an associated value 430 and a score 440.

The UPI attributes 420 may include information concerning user background and interests such as, for example, geographic information, age or age group, topics of interest, reading level, income and other demographics suited for targeting advertisements. A value 430 associated with a UPI attribute 420 may be quantitative (a discrete or continuous value, e.g.,

5 Age=58 years; Annual income=\$55,000; City of residence=San Francisco, CA) or qualitative (in set or not in set, e.g., Salary \$50,000-\$100,000?=Yes; U.S. Resident?=Yes). Each UPI attribute 420 and value 430 may have an associated score 440 related to the probability of the attribute value being correct. Examples of UPI attributes may include one or more of the following:

- 10 - the content (e.g., words, Anchortext, etc.) of Websites that the user has visited (or visited in a certain time period);
- demographic information;
- geographic information;
- psychographic information;
- previous queries (and/or associated information) that the user has made;
- 15 - information about previous advertisements that the user has been shown, has selected, and/or has made purchases after viewing;
- information about documents (e.g., word processor) viewed/requested, and/or edited by the user;
- user interests;
- 20 - explicit or implicit feedback regarding the personalized results from the user (e.g., selecting a result, not selecting a result, the amount of time spent on a result, etc.);
- browsing activity; and
- previous purchasing behavior.

25 Such information may be stored on a per individual basis (though it might not include personally identifiable information for reasons of privacy), or aggregated in various ways among various sets of individuals. Such information may be combined to obtain composite profiles.

#### 30 § 4.2.2 EXEMPLARY APPLICATIONS FOR USING UPI

There are many applications for using user profile information (UPI). Three exemplary applications – enhanced ad targeting, ad creative selection and generation, and resolving query ambiguity -- are described below.



### § 4.2.2.1 AD SCORING USING UPI

It may be desirable to place an ad on a hosting site or page where ad's targeted UPI (and/or the UPI of the ad's landing page) matches the document UPI (e.g., the average UPI of users that have requested the document) and/or to serve the ad to a user whose UPI matches the target UPI of the ad (and/or the UPI of the ad's landing page).

With enhanced ad targeting using UPI, a score for each of a plurality of ads may be determined using at least some of the UPI of the user, the UPI of a document, the UPI of an ad landing page, and/or ad targeting UPI. At least one ad may be rank ordered, filtered, and/or selected from the plurality of ads using at least the determined scores.

For example, an ad score may be a function of a UPI match value. Referring to Figure 5, such a UPI match value may be a function of one or more of: (i) a match value of UPI information 514 associated with a user (or user group) 512 and UPI information 524 associated with a document (e.g., a Web page) requested (or visited) by the user 512 (Match Value 1); (ii) a match value of UPI information 514 associated with a user (or user group) 512 and UPI information 534 associated with a landing page 532 of an ad under consideration (Match Value 2); (iii) a match value of UPI information 524 associated with a document 522 requested (or visited) by a user and UPI information 534 associated with a landing page 532 of an ad under consideration (Match Value 3); (iv) a match value of UPI information 514 associated with a user (or user group) 512 and UPI ad targeting information 544 associated with an ad 542 under consideration (Match Value 4); and (v) a match value of UPI information 524 associated with a document 522 requested or being visited by a user and UPI ad targeting information 544 associated with an ad 542 under consideration (Match Value 5).

Thus, in one embodiment of the present invention, an overall match may be defined as:

$$\begin{aligned} \text{UPI MATCH} = & a * \text{MATCH VALUE 1} + b * \text{MATCH VALUE 2} + \\ & c * \text{MATCH VALUE 3} + d * \text{MATCH VALUE 4} + \\ & e * \text{MATCH VALUE 5} \end{aligned}$$

where a, b, c, d, and e are constants (e.g., a=0.025, b=0.275, c=0.3, d=0.2, and e=0.2). Note that "a" may be set to zero since a match between a user UPI and a document UPI may be independent of how well either one matches an ad. Other functions, including polynomial or exponential functions, may be used instead.

Generally, for a good match, the user's UPI should match both the hosting page UPI, and perhaps even more importantly, match the UPI of the ad landing page. The match between two profiles can be computed using standard Information Retrieval techniques for matching two term vectors, such as vector space matching (See, e.g., the articles: G. Salton and C. Buckley, "Term-Weighting Approaches in Automatic Text Retrieval," Information Processing and Management, 24(5), pp. 513-523 (1988); and Gerard Salton, A. Wong, C. S. Yang, "A Vector Space Model for Automatic Indexing," Communications of the ACM, 18(11), pp. 613-620 (1975).).

Broad attributes such as geography, topic, user age range, language, etc. can be computed for documents and users using, for example, machine learning classifiers. Also, that such broad attributes can be used jointly with more narrow attributes, such as words and phrases, in matching.

Naturally, the score of an ad can be a function of other factors in addition to UPI Match, such as, for example, its relevancy to a search query or to content of a document, an amount an advertiser will pay or is willing to pay for a given result (e.g., impression, selection, conversion, etc.), a measure of the ad's performance (e.g., click-through rate, conversion rate, user rating, etc.), a measure of the advertiser's quality, etc. Moreover, different intermediate ad scores may be used for different purposes (e.g., relevancy, position, relative rendering attribute, etc.)

#### 20 § 4.2.2.2 AD CREATIVE SELECTION USING UPI

With ad creative selection and/or generation using UPI, ad creatives may be tailored to UPI (e.g., of the user, the document, etc.)

#### 25 § 4.2.2.3 QUERY ABIGUITY RESOLUTION USING UPI

Besides selecting an appropriate ad or ad creative, UPI may be used to disambiguate a search query including an ambiguous search term. For example, a user submitting the search query "jaguar" might want either (A) information about Jaguar cars, (B) information about the animal, (C) information about the Apple Jaguar operating system, or (D) information about the Jacksonville Jaguar NFL football team. UPI of the user could be used to help disambiguate the "jaguar" search term. In this example, information from the user's previous browsing activity may help disambiguate the ambiguous query "jaguar", or prevent the display of advertisements of little or no interest to the user. For example, if the user had been recently querying for "apple

computer" and "operating systems", and subsequently submitted the search query "jaguar", the user's previous query history could be used to infer that the query was more likely referring to the operating system for Apple computers and not to the car, animal, or NFL team.

Consequently, ads could be better targeted to the user by harnessing such information.

5           These foregoing applications for using UPI, as well as others, presume that UPI is available. If, however, this is not the case, such UPI needs to be generated, and/or updated. Exemplary methods for determining and/or updating UPI are described in § 4.2.3 below.

### § 4.2.3 DETERMINING AND/OR UPDATING UPI

10

Recall that user profile information (UPI) can be determined using information provided by users when they subscribe to a service. However, in many cases, users will not volunteer information about themselves, or are simply not asked to volunteer such information. Further, even if users do volunteer such information, it may be incomplete (e.g., because the information given at the time of subscription may be limited to what is needed for the service and hence not comprehensive), it may be intentionally or unintentionally inaccurate, it may become stale, etc.

15           Similarly, UPI information may not be available for document which the user has requested or is visiting.

Although ads can have advertisers supply UPI information, doing so may be a burden to advertisers.

20           As described below, UPI for a user, a document, etc. can be determined (or updated or extended) even when no explicit information is given to the system. In the following, an initial (or baseline) UPI may be determined during a bootstrap phase. Such an initial UPI may use information inherent to the user, or to the document, or to the ad with which the UPI will be associated. An initial UPI may include some expressly entered UPI information, though it doesn't need to.

25           It may suffice to use such an initial or baseline UPI for applications such as ad scoring, ad creative selection, query ambiguity resolution. However, it may be desired to supplement the initial or baseline UPI during an expansion and/or reinforcement phase. Referring back to Figure 4 for example, UPI 410 (e.g., of a user, or document, or ad, or ad landing page) may be expanded by adding values 430 to attributes 420 that either didn't exist previously, or that didn't have a value. Alternatively, or in addition, scores of 440 of a value 430 of an attribute 420 of the UPI 410 may be revised or reinforced. Thus, for example, the probability of attribute=sex

having a value=male may increase or decrease given additional information. Information used to expand or supplement an initial UPI may be thought of as inferred information.

Figure 6 is a flow diagram of an exemplary method 600 that may be used to determine document UPI in a manner consistent with the present invention. An initial UPI for the document is determined. (Block 610) This may be done using information inherent to the document, such as the document content. An exemplary method for performing this act is described below with reference to Figure 8. Further UPI information for the document is inferred. (Block 620) This may be done using information inferred from visiting users, linked documents, etc. Then, a new (e.g., expanded and/or reinforced) document UPI is determined using the initial document UPI and the inferred document UPI. (Block 630) Acts 620 and 630 may be performed one or more times before the method 600 is left. (Node 640). If acts 620 and 630 are repeated, as indicated by the phantom branch-back line, the new document UPI may be determined using the previous document UPI and more inferred document UPI.

Figure 7 is a flow diagram of an exemplary method 700 that may be used to determine user (as an individual or a group) UPI in a manner consistent with the present invention. An initial UPI for the user is determined. (Block 710) This may be done using information inherent to the user, such as past search queries submitted by the user. An exemplary method for performing this act is described below with reference to Figure 9. Further UPI information for the user is inferred. (Block 720) This may be done using information inferred from search results, selected search results, etc. Then, a new (e.g., expanded and/or reinforced) user UPI is determined using the initial user UPI and the inferred user UPI. (Block 730) Acts 720 and 730 may be performed one or more times before the method 700 is left. (Node 740). If acts 720 and 730 are repeated, as indicated by the phantom branch back-line, the new user UPI may be determined using the previous user UPI and more inferred user UPI.

Recall from block 610 of Figure 6 that initial document UPI may be determined. Figure 8 is a flow diagram of an exemplary method 800 that may be used to determine an initial or baseline document UPI in a manner consistent with the present invention. As shown, document UPI attribute values may be populated (e.g., using content of the document, meta information of the document, etc.). (Block 810)

Recall from block 710 of Figure 7 that initial user UPI may be determined. Figure 9 is a flow diagram of an exemplary method 900 that may be used to determine an initial or baseline user UPI in a manner consistent with the present invention. As shown, user UPI attribute values may be populated (e.g., using past query information of the user, information expressly entered by the user, etc.). (Block 910)

#### § 4.2.4 EXAMPLES OF HOW USER INFORMATION IS OBTAINED

5           There are many alternative ways to obtain user information. For example, a score 440 for an attribute 420 and value 430 can be determined with a machine learning classifier which predicts values 430 of the UPI attributes 420 in the profile using words in queries deployed previously. For example, given the keywords related to "women's health" in previous search queries, the classifier may infer that the user is a woman with probability 0.8. Further, given 10 that Japanese words were used in previous search queries, the classifier may infer that the user is Japanese with probability 0.9, etc. The scores 440 of two UPIs may affect whether or not they match. Attribute 420 may be "San Francisco" in one UPI and "San Jose" in another. Although these raw attributes do not match, a machine learning classifier may assign a geographic category = "California" in both cases and the generalized profile attributes will match. Thus, the 15 machine learning classifier may be used to classify users and documents into broad categories relating to geography, topic, ethnicity, or reading level.

          Recall from both block 620 of Figure 6 and 720 of Figure 7 that UPI information may be inferred. User UPI may be inferred from UPI of other users and/or documents associated in some way with the user. Similarly, document UPI may be inferred from UPI of other 20 documents, and/or users associated in some way with the document. Figure 10 illustrates how users and documents can be associated.

          As shown in Figure 10, user/document association operations 1010 may generate information 1070 associating users and/or documents (user-to-user, user-to-document, document-to-document, and/or document-to-user (which may be the same as a user-to-document 25 association in an undirected graph such as the one shown)) using one or more of user information 1020 of one or more users, document information 1030 of one or more documents, document link (e.g., Web topology) information 1040, search results for user queries 1050 (which may be provided as user information 1020 instead), and user document selections 1060 (which may be provided as user information 1030 instead).

30           In one exemplary embodiment of the present invention, the association information 1070 may be a graph in which users and documents are represented as nodes 1072 and 1076, respectively. Figure 11 is a flow diagram of an exemplary method 1100 that may be used to associate users and/or documents in a manner consistent with the present invention. As shown, nodes may be defined for each user and document. (Block 1110) For each of the user nodes 35 1072, edges 1074 (which indicate an association) may be drawn between the user node and

document nodes for the top Web pages that were returned by a search engine in response to search queries that the user submitted. (In a variant, the edges 1074 could be drawn only to Web pages that the user selected (e.g., clicked on)). Additionally, edges 1078 may be drawn between pairs of documents that have links (e.g., hyperlinks) between them. (Block 1120) Although not shown, user-to-user associations may also be generated. For example, edges may be added between users that have visited one or more of the same documents.

Referring back to Figure 7, given this graph 1070, new UPI for a user may be determined using their initial UPI and, as inferred user UPI, the UPIs of documents corresponding to nodes 1076 with which the user node 1072 is linked. Similarly, referring back to Figure 6, given this graph 1070, new UPI for a document may be determine using its initial UPI and, as inferred document UPI, the UPIs of users and documents corresponding to nodes 1072 and 1076, respectively, with which the document node 1076 is linked.

Referring back to blocks 630 and 730 of Figures 6 and 7 respectively, in one embodiment of the present invention:

$$\text{updated\_UPI} = \text{initial\_UPI} * \alpha + \text{neighborhood\_UPI} * \beta$$

where alpha and beta are constants (e.g., alpha = 0.7 beta = 0.3) and neighborhood\_UPI is the average of the UPIs of the neighboring nodes in the graph. Other functions for updating UPI are possible. This process may be repeated for a number (e.g., 50) of iterations.

To determine an average UPI, the values 430 of individual attributes 420 may be averaged, perhaps weighted by scores 440.

More distant nodes (e.g., two or more edges away) may also be considered, but should be weighted less.

Although both user UPI and document UPI may be updated, it is possible to update only one or the other.

In alternative embodiments, a user node 1072 may represent an aggregation of users. Similarly, a document node 1076 may represent an aggregation of documents (e.g., a Website containing a number of Web pages).

In alternative embodiments, graph edges can be assigned association weights. Thus, for example, an edge 1074 from a user node 1072 to a document node 1076 may be provided with a larger weight if the user selected the document from a search result list, than if the document was merely included on a search results list returned in response to a user search query.

### § 4.2.5 EXEMPLARY APPARATUS

Figure 12 is high-level block diagram of a machine 1200 that may effect one or more of the operations discussed above (e.g., those performed by an end user system or client device, those performed by a content server, those performed by a search engine, or those performed by an ad server). The machine 1200 basically includes one or more processors 1210, one or more input/output interface units 1230, one or more storage devices 1220, and one or more system buses and/or networks 1240 for facilitating the communication of information among the coupled elements. One or more input devices 1232 and one or more output devices 1234 may be coupled with the one or more input/output interfaces 1230.

The one or more processors 1210 may execute machine-executable instructions (e.g., C or C++ running on the Solaris operating system available from Sun Microsystems Inc. of Palo Alto, California or the Linux operating system widely available from a number of vendors such as Red Hat, Inc. of Durham, North Carolina) to effect one or more aspects of the present invention. At least a portion of the machine executable instructions may be stored (temporarily or more permanently) on the one or more storage devices 1220 and/or may be received from an external source via one or more input interface units 1230.

In one embodiment, the machine 1200 may be one or more conventional personal computers. In this case, the processing units 1210 may be one or more microprocessors. The bus 1240 may include a system bus. The storage devices 1220 may include system memory, such as read only memory (ROM) and/or random access memory (RAM). The storage devices 1220 may also include a hard disk drive for reading from and writing to a hard disk, a magnetic disk drive for reading from or writing to a (e.g., removable) magnetic disk, and an optical disk drive for reading from or writing to a removable (magneto-) optical disk such as a compact disk or other (magneto-) optical media.

A user may enter commands and information into the personal computer through input devices 1232, such as a keyboard and pointing device (e.g., a mouse) for example. Other input devices such as a microphone, a joystick, a game pad, a satellite dish, a scanner, or the like, may also (or alternatively) be included. These and other input devices are often connected to the processing unit(s) 1210 through an appropriate interface 1230 coupled to the system bus 1240. The output devices 1234 may include a monitor or other type of display device, which may also be connected to the system bus 1240 via an appropriate interface. In addition to (or instead of) the monitor, the personal computer may include other (peripheral) output devices (not shown), such as speakers and printers for example.

### § 4.3 REFINEMENTS AND ALTERNATIVES

#### § 4.3.1 DETERMINATION TECHNIQUES

5

Matching (or similarity) determination can be performed in a number of ways. For example, one or more of the following similarity determination techniques may be used: (a) vector-based (as described below for example); (b) rule-based (as described below for example); (c) probabilistic reasoning to infer a probability or likelihood of match; and (d) fuzzy logic matching. Other similarity determination techniques may be used under the present invention as well.

As just alluded to, one way of determining similarity is to form UPI vectors. Each of a plurality of advertisements may also have a representative targeting UPI vector. Note that such UPI vectors may also include additional information determined from historical data. For example, inferences about which demographic groups tend to click on an ad may be determined using data mining techniques. The results of this (e.g., off-line) analysis may then be used as one or more attributes in a UPI targeting vector of the ads. UPI vectors can be matched using a scoring function such as the cosine distance between the vectors, a hamming distance, and/or any one of a variety of other vector distance measures. Advertisements can then be ranked using (e.g., according to) the scores generated by such a function.

As also indicated above, another technique for matching UPI involves applying a set of rules and/or functions that define a similarity of pairs of one or more of user, document, ad and ad landing page UPIs. Such a rule and/or formula-based system can use arbitrary combinations of logical rules (e.g., with weights attached) to give weighted scores.

As stated above, the vector-based techniques and the rules and/or formula-based techniques may be used in concert to generate a match determination.

In one embodiment of the present invention, an attribute frequency - inverse document (or user) frequency product (TF-IDF) measure is determined using UPIs to generate a similarity score.

30

#### § 4.3.1.1 REFINEMENTS TO SIMILARITY DETERMINATION TECHNIQUES

To keep the process of matching UPIs efficient, attribute selection, attribute generalization, and/or attribute aggregation techniques may be used to reduce the sizes of the

35



vectors being compared. For example, attribute selection may be used to reduce the number, and/or size of the attributes used as part of the matching process between UPIs. Such attribute selection techniques may include keeping only some number of attributes having a high score (in absolute and/or relative terms) (e.g., probability or importance weight) in the vector  
5 representation of the UPIs. Another attribute selection technique may include using a statistical measure, such as “mutual Information”, Chi-squared fit, or correlation for example, to determine which attributes are more indicative than others for generating one or more matches that are likely to perform well (e.g., be clicked on). Yet another attribute selection technique is simply hand-selecting those attributes believed to be most useful. Two or more of these or other  
10 attribute selection techniques may be used in concert.

Similarly, ad UPI vectors may be pre-filtered using some set of initial criteria (for example, matching certain features exactly) so that only a small subset of the ad UPI vectors remaining will need to be ranked with respect to the user and of document UPIs. This technique illustrates a combined rule-based and vector distance similarity determination technique. Using  
15 this technique has the added advantage of helping to increase the efficiency of the overall matching process between the UPIs.

For example, a first set of one or more ads generally relevant to a user and/or a document may be determined first. This first set of one or more ads may be processed further using UPI information operations to determine a final set of ads (e.g., ordered and/or filtered based on a  
20 second similarity score determined using UPI). Such an embodiment may be useful when (part of) the UPI is stored on the client. For example, an ad server may send the top N (e.g., N=100) ads for a user request (e.g., a search query or a document request) to the client. The client can then reorder these served ads based on the local user information.

### 25 § 4.3.2 EXAMPLES OF HOW AND WHERE UPI IS STORED

In one embodiment of the present invention, UPI is stored as n-grams (sequences of words of length 1 to some maximum, e.g., 3) with an associated source (e.g., the n-gram is from a Web page viewed, a query, etc.), URLs, or specific values for different kinds of information  
30 (e.g., the user is using a particular web browser, or the user is located at a specific latitude and longitude, etc.).

The storage and use of UPI may take various different forms, including (a) client-side storage (e.g., in the form of a browser cookie, local file, hidden form fields, or URL encoding), (b) server-side storage (e.g., a database of records, flat files, or proprietary structures for storing

and retrieving profile/behavior information), and/or (c) third party storage. Thus, the UPI need not reside in a repository on the server-side, but may actually be stored with the client and/or a third party and sent to the information server with the user's request (e.g., search query or document request). The UPI can be sent, for example, in a web browser cookie.

5 Referring back to the exemplary network environment of Figure 3, UPI may be stored in one or more client devices 320, one or more content servers 330, one or more search engines 340, one or more ad servers 350, and/or one or more user information servers 360. In any event, user information may be maintained by using some data management scheme (e.g., database, flat files, proprietary data management system, web browser cookies, etc.). Different items of  
10 user information (e.g., different user features) may be stored on different devices.

#### § 4.3.2.1 REFINEMENTS TO INFORMATION STORAGE

15 To achieve a practical system, it may become desirable or necessary to reduce the volume of, and/or to order, UPI used in targeting advertisements. The present invention permits such reduction with minimal loss of information and accuracy. Further, since it may be desirable or necessary to constrain latency in the serving of ads, and consequently in the retrieval of UPI, an ad server may need to operate in some cases with limited or no UPI. Prioritization, in  
20 a manner consistent with the present invention, allows an ad server to focus on valuable (e.g., the most valuable) information available during the targeting process. Exemplary techniques for data reduction and processing prioritization are described below.

In one embodiment of the present invention, each of at least some UPI attributes has an associated score (importance weight, probability, etc.). The importance weights for multiple  
25 UPI attributes may be grouped together or aggregated in some manner.

Initially, the importance weight assigned to a UPI attribute may be determined using, perhaps among other things, a type of the information. For example search queries may have a higher initial weight than n-grams extracted from Web pages that a user has viewed.

The importance weights for UPI attributes may be updated based on actions of a  
30 particular user or of groups of users. For example, Figure 9 of U.S. Patent Application Serial No. 10/452,791, titled "SERVING ADVERTISEMENTS USING USER REQUEST INFORMATION AND USER INFORMATION," filed on June 2, 2003 and listing Steve Lawrence, Mehran Sahami and Amit Singhal as inventors (incorporated herein by reference and referred to as "the '791 application") illustrates an exemplary method that may be used to  
35 manage ad UPI or user UPI in a manner consistent with the present invention. In that exemplary

method, for example, if the user selects an ad, the importance weights of user UPI attributes that led to this ad being recommended may be increased. The amount(s) (absolute or relative) by which the importance weight(s) are increased may depend, at least in part, on how much each UPI attribute contributed to the ad recommendation score. Similarly, the importance weights of the relevant UPI attributes may be reduced when the user does not select the ad. In one embodiment of the present invention, the weight reduction factor is smaller than the weight increase factor.

In one embodiment of the present invention, importance weights may be optionally adjusted using (e.g., combined with) a global weight based on how often (in absolute and/or relative terms) the UPI attribute has resulted in a good recommendation across all users. Alternatively, or in addition, importance weights may be optionally adjusted using (e.g., combined with) a community weight based on how often (in absolute and/or relative terms) the UPI attribute has resulted in a good recommendation across a group of (e.g., all similar) users.

In one embodiment of the present invention, if a UPI reaches a particular size, it may be reduced by deleting less important items of information using, at least, the importance weights. The storage requirements of the UPI attributes may also be considered. In this way, a UPI attribute that requires a lot of storage but that has a low importance weight would be a prime candidate for deletion.

One embodiment of the present invention can order (and even store) the individual UPI attributes using their importance weights. In such an embodiment, requests for a given user UPI may be served with truncated UPI including only the most important (in terms of importance weights) attributes, or may be served in a piecemeal manner until enough information is obtained, and/or until further processing is undesirable or not possible. Thus, for example, processing using UPI attributes may be done in the order of their importance weights. This permits processing to be terminated early (e.g., for efficiency, to meet some latency constraint and ensure that processing is completed within a specific time limit, etc.).

As new UPI attributes are added, the importance weights of existing attributes may be changed. Consequently, the order of UPI attributes (e.g., for a given user) may differ from the importance weight order. Therefore, the present invention may be used to reorder, periodically, the UPI attributes using (e.g., according to) their respective importance weights.

In some cases, transfer or processing of the UPI may be performed in accordance with an approximate order, before the UPI (e.g., for a given user) is reordered using (e.g., in accordance with) importance weight. Note that in at least some instances and/or some embodiments, it may not be necessary to physically store UPI attributes in the order of their respective importance

weights for such information to be efficiently transferred and/or processed in that order. Indeed, in at least some instances and/or some embodiments, it may not be necessary to transfer and/or process UPI attributes in importance order. For example, if conditions permit, all information may be transferred and/or processed.

5

#### § 4.4 CONCLUSIONS

As can be appreciated from the foregoing disclosure, the present invention can be used to improve content-targeted ad systems, as well as keyword-targeted ad systems. User profile  
10 information of users, documents, ads and/or ad landing pages can be accepted and/or inferred.

**WHAT IS CLAIMED IS:**

- 1 1. A method for determining user profile information for a user, the method comprising:
  - 2 a) determining initial user profile information for the user;
  - 3 b) inferring user profile information for the user; and
  - 4 c) determining the user profile information for the user using both the initial user profile
  - 5 information and the inferred user profile information.
  
- 1 2. The method of claim 1 wherein the act of determining an initial user profile information for
- 2 the user uses past search queries submitted by the user.
  
- 1 3. The method of claim 1 wherein the act of determining an initial user profile information for
- 2 the user uses past document selections by the user.
  
- 1 4. The method of claim 1 wherein the act of determining initial user profile information for the
- 2 user uses (i) past search queries submitted by the user, and (ii) past document selections by the
- 3 user.
  
- 1 5. The method of claim 1 wherein the initial user profile includes a plurality of attributes, each
- 2 of the plurality of attributes having a value and a score.
  
- 1 6. The method of claim 5 wherein the score indicates a likelihood that the value of the attribute
- 2 is correct.
  
- 1 7. The method of claim 1 wherein the act of inferring user profile information for the user
- 2 includes
  - 3 i) defining a node for each of a number of documents and the user,
  - 4 ii) adding edges between nodes if there is an association between the nodes to
  - 5 define a graph, and
  - 6 iii) inferring user profile information for the user using a topology of the graph
  - 7 and user profile information of other documents.

1 8. The method of claim 7 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was returned in a search results page to a search query  
3 from the user corresponding to the second node.

1 9. The method of claim 7 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was selected by the user corresponding to the second  
3 node.

1 10. The method of claim 7 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node is linked with a document corresponding to the second  
3 node.

1 11. The method of claim 7 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was visited by a set of users that have visited another  
3 document corresponding to the second node.

1 12. The method of claim 7 wherein an edge is added between first and second nodes if a user  
2 corresponding to the first node visited a set of one or more documents also visited by another  
3 user corresponding to the second node.

1 13. The method of claim 7 wherein the act of inferring user profile information for the user  
2 using a topology of the graph includes  
3 i) multiplying the initial user profile information of the user by a first value to generate a  
4 first product;  
5 ii) multiplying user profile information of neighboring graph nodes by a second value to  
6 generate a second product; and  
7 iii) adding the first product and the second product.

1 14. A method for determining user profile information for a document, the method comprising:  
2 a) determining initial user profile information for the document;  
3 b) inferring user profile information for the document; and  
4 c) determining the user profile information for the document using both the initial user  
5 profile information and the inferred user profile information.

- 1 15. The method of claim 14 wherein the act of determining an initial user profile information  
2 for the document uses content information from the document.
- 1 16. The method of claim 14 wherein the act of determining initial user profile information for  
2 the document uses document meta information.
- 1 17. The method of claim 14 wherein the act of determining initial user profile information for  
2 the document uses (i) content information from the document, and (ii) document meta  
3 information.
- 1 18. The method of claim 14 wherein the initial user profile information includes a plurality of  
2 attributes, each of the plurality of attributes having a value and a score.
- 1 19. The method of claim 18 wherein the score indicates a likelihood that the value of the  
2 attribute is correct.
- 1 20. The method of claim 14 wherein the act of inferring user profile information for the  
2 document includes  
3 i) defining a node for each of a number of documents and for each of a number  
4 of users,  
5 ii) adding edges between nodes if there is an association between the nodes to  
6 define a graph, and  
7 iii) inferring user profile information for the document using a topology of the  
8 graph and user profile information of users and of other documents.
- 1 21. The method of claim 20 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was returned in a search results page to a search query  
3 from the user corresponding to the second node.
- 1 22. The method of claim 20 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was selected by the user corresponding to the second  
3 node.

1 23. The method of claim 20 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node is linked with a document corresponding to the second  
3 node.

1 24. The method of claim 20 wherein an edge is added between first and second nodes if a  
2 document corresponding to the first node was visited by a set of users that have visited another  
3 document corresponding to the second node.

1 25. The method of claim 20 wherein an edge is added between first and second nodes if a user  
2 corresponding to the first node visited a set of one or more documents also visited by another  
3 user corresponding to the second node.

1 26. The method of claim 20 wherein the act of inferring user profile information for the  
2 document using a topology of the graph includes  
3 i) multiplying the initial user profile information of the document by a first value to  
4 generate a first product;  
5 ii) multiplying user profile information of neighboring graph nodes by a second value to  
6 generate a second product; and  
7 iii) adding the first product and the second product.

1 27. A method for determining a match used for scoring an ad, the method comprising:  
2 a) determining a first match value using (A) at least one of user profile information of an  
3 ad landing page of the ad and user profile information used for targeting the ad, and (B)  
4 user profile information of a user to which the ad will be rendered;  
5 b) determining a second match value using (A) at least one of user profile information of  
6 an ad landing page of the ad and user profile information used for targeting the ad, and  
7 (B) user profile information of a document with which the ad will be served; and  
8 c) determining the match used for scoring the ad using the first match value and the  
9 second match value.

1 28. The method of claim 27 wherein at least some of the user profile information of the ad  
2 landing page of the ad was inferred.



- 1 29. The method of claim 27 wherein at least some of the user profile information used for  
2 targeting of the ad was inferred.
- 1 30. The method of claim 27 wherein at least some of the user profile information of the user  
2 was inferred.
- 1 31. The method of claim 27 wherein at least some of the user profile information of the  
2 document was inferred.
- 1 32. The method of claim 27 wherein the user profile information includes  
2 - at least one broad attribute selected from a set of broad attributes including (A) a  
3 geographic area, (B) a topic, (C) a user age range, (D) a language, and  
4 - at least one narrow attribute selected from a set of narrow attributes including (A)  
5 words, and (B) phrases.
- 1 33. Apparatus for determining user profile information for a user, the apparatus comprising:  
2 a) means for determining initial user profile information for the user;  
3 b) means for inferring user profile information for the user; and  
4 c) means for determining the user profile information for the user using both the initial  
5 user profile information and the inferred user profile information.
- 1 34. The apparatus of claim 33 wherein the means for determining an initial user profile  
2 information for the user use past search queries submitted by the user.
- 1 35. The apparatus of claim 33 wherein the means for determining an initial user profile  
2 information for the user use past document selections by the user.
- 1 36. The apparatus of claim 33 wherein the means for determining initial user profile  
2 information for the user use (i) past search queries submitted by the user, and (ii) past document  
3 selections by the user.
- 1 37. The apparatus of claim 33 wherein the initial user profile includes a plurality of attributes,  
2 each of the plurality of attributes having a value and a score.

1 38. The apparatus of claim 37 wherein the score indicates a likelihood that the value of the  
2 attribute is correct.

1 39. The apparatus of claim 37 wherein the means for inferring user profile information for the  
2 user include means for

- 3 i) defining a node for each of a number of documents and the user,
- 4 ii) adding edges between nodes if there is an association between the nodes to  
5 define a graph, and
- 6 iii) inferring user profile information for the user using a topology of the graph  
7 and user profile information of other documents.

1 40. The apparatus of claim 39 wherein the means for adding edges adds an edge between first  
2 and second nodes if a document corresponding to the first node was returned in a search results  
3 page to a search query from the user corresponding to the second node.

1 41. The apparatus of claim 39 wherein the means for adding edges adds an edge between first  
2 and second nodes if a document corresponding to the first node was selected by the user  
3 corresponding to the second node.

1 42. The apparatus of claim 39 wherein the means for adding edges adds an edge between first  
2 and second nodes if a document corresponding to the first node is linked with a document  
3 corresponding to the second node.

1 43. The apparatus of claim 39 wherein the means for adding edges adds an edge between first  
2 and second nodes if a document corresponding to the first node was visited by a set of users that  
3 have visited another document corresponding to the second node.

1 44. The apparatus of claim 39 wherein the means for adding edges adds an edge between first  
2 and second nodes if a user corresponding to the first node visited a set of one or more documents  
3 also visited by another user corresponding to the second node.

1 45. The apparatus of claim 39 wherein the means for inferring user profile information for the  
2 user using a topology of the graph include means for

- 3 i) multiplying the initial user profile information of the user by a first value to generate a  
4 first product;  
5 ii) multiplying user profile information of neighboring graph nodes by a second value to  
6 generate a second product; and  
7 iii) adding the first product and the second product.

1 46. Apparatus for determining user profile information for a document, the apparatus  
2 comprising:

- 3 a) means for determining initial user profile information for the document;  
4 b) means for inferring user profile information for the document; and  
5 c) means for determining the user profile information for the document using both the  
6 initial user profile information and the inferred user profile information.

1 47. The apparatus of claim 46 wherein the means for determining an initial user profile  
2 information for the document use content information from the document.

1 48. The apparatus of claim 46 wherein the means for determining initial user profile  
2 information for the document use document meta information.

1 49. The apparatus of claim 46 wherein the means for determining initial user profile  
2 information for the document use (i) content information from the document, and (ii) document  
3 meta information.

1 50. The apparatus of claim 46 wherein the initial user profile information includes a plurality of  
2 attributes, each of the plurality of attributes having a value and a score.

1 51. The apparatus of claim 50 wherein the score indicates a likelihood that the value of the  
2 attribute is correct.

1 52. The apparatus of claim 46 wherein the means for inferring user profile information for the  
2 document include means for

- 3 i) defining a node for each of a number of documents and for each of a number  
4 of users,

- 5           ii) adding edges between nodes if there is an association between the nodes to  
6           define a graph, and  
7           iii) inferring user profile information for the document using a topology of the  
8           graph and user profile information of users and of other documents.

1   53. The apparatus of claim 52 wherein the means for adding edges adds an edge between first  
2   and second nodes if a document corresponding to the first node was returned in a search results  
3   page to a search query from the user corresponding to the second node.

1   54. The apparatus of claim 52 wherein the means for adding edges adds an edge between first  
2   and second nodes if a document corresponding to the first node was selected by the user  
3   corresponding to the second node.

1   55. The apparatus of claim 52 wherein the means for adding edges adds an edge between first  
2   and second nodes if a document corresponding to the first node is linked with a document  
3   corresponding to the second node.

1   56. The apparatus of claim 52 wherein the means for adding edges adds an edge between first  
2   and second nodes if a document corresponding to the first node was visited by a set of users that  
3   have visited another document corresponding to the second node.

1   57. The apparatus of claim 52 wherein the means for adding edges adds an edge between first  
2   and second nodes if a user corresponding to the first node visited a set of one or more documents  
3   also visited by another user corresponding to the second node.

1   58. The apparatus of claim 52 wherein the means for inferring user profile information for the  
2   document using a topology of the graph include means for  
3           i) multiplying the initial user profile information of the document by a first value to  
4           generate a first product;  
5           ii) multiplying user profile information of neighboring graph nodes by a second value to  
6           generate a second product; and  
7           iii) adding the first product and the second product.

1   59. Apparatus for determining a match used for scoring an ad, the apparatus comprising:

- 2 a) means for determining a first match value using (A) at least one of user profile  
3 information of an ad landing page of the ad and user profile information used for  
4 targeting the ad, and (B) user profile information of a user to which the ad will be  
5 rendered;
- 6 b) means for determining a second match value using (A) at least one of user profile  
7 information of an ad landing page of the ad and user profile information used for  
8 targeting the ad, and (B) user profile information of a document with which the ad will  
9 be served; and
- 10 c) means for determining the match used for scoring the ad using the first match value  
11 and the second match value.

1 60. The apparatus of claim 59 wherein at least some of the user profile information of the ad  
2 landing page of the ad was inferred.

1 61. The apparatus of claim 59 wherein at least some of the user profile information used for  
2 targeting of the ad was inferred.

1 62. The apparatus of claim 59 wherein at least some of the user profile information of the user  
2 was inferred.

1 63. The apparatus of claim 59 wherein at least some of the user profile information of the  
2 document was inferred.

1 64. The apparatus of claim 59 wherein the user profile information includes  
2 - at least one broad attribute selected from a set of broad attributes including (A) a  
3 geographic area, (B) a topic, (C) a user age range, (D) a language, and  
4 - at least one narrow attribute selected from a set of narrow attributes including (A)  
5 words, and (B) phrases.

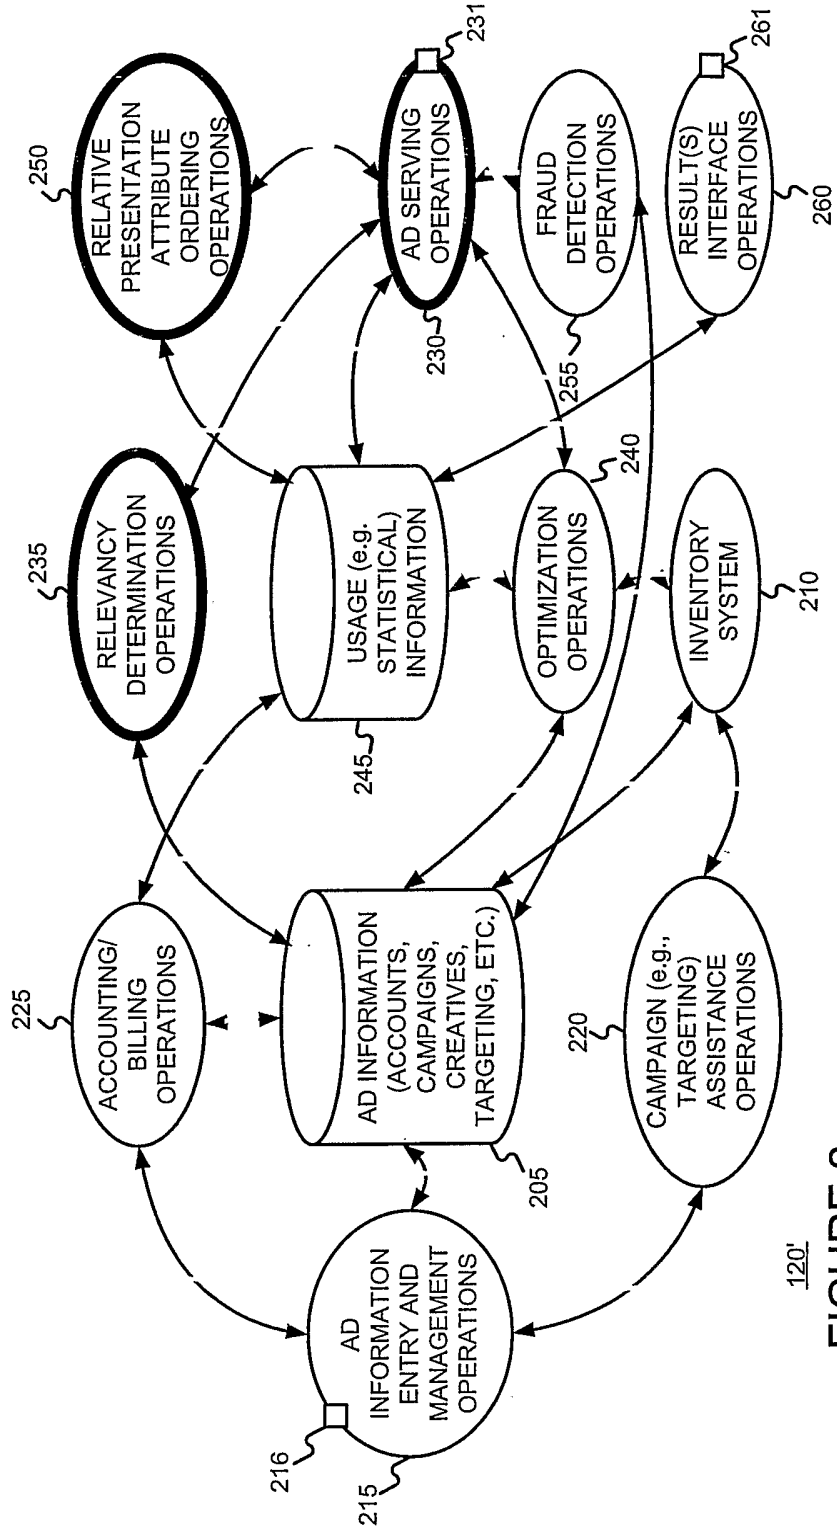

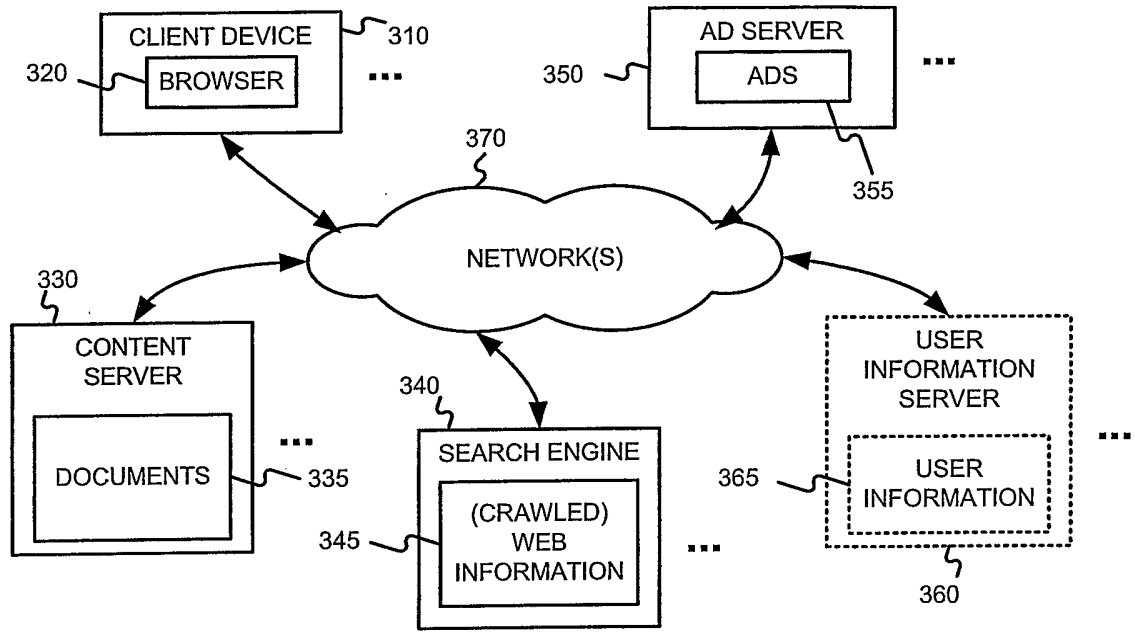


FIGURE 2



300

FIGURE 3




FIGURE 4




FIGURE 5



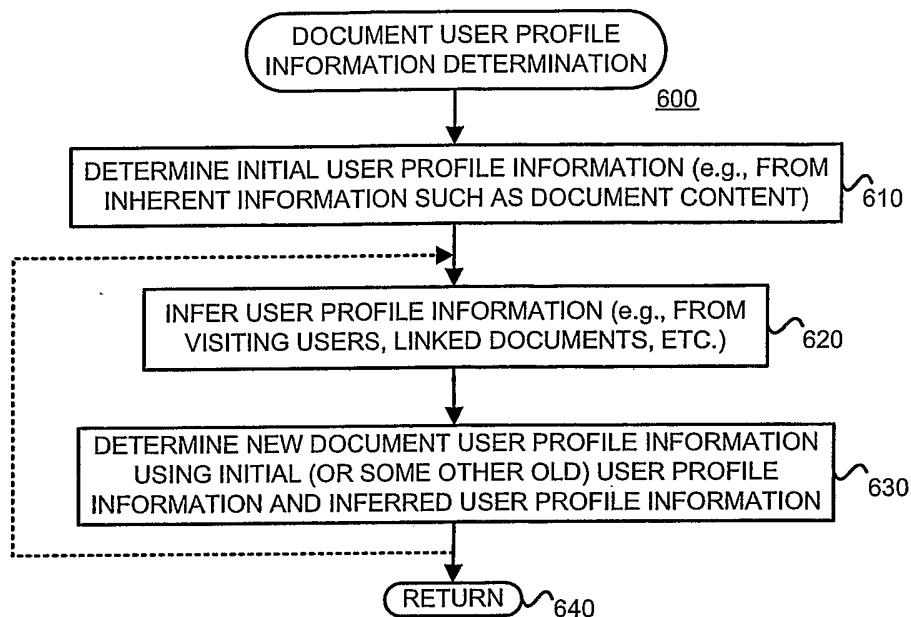


FIGURE 6

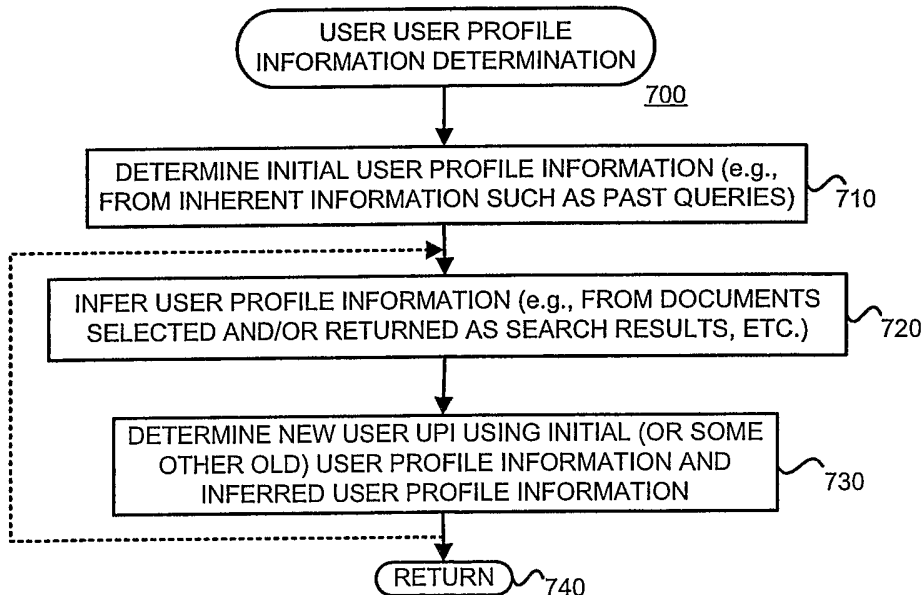


FIGURE 7

5/7




FIGURE 8

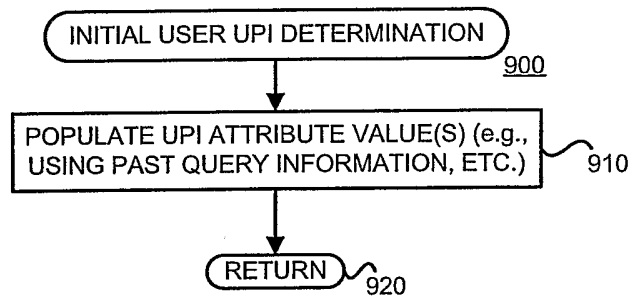


FIGURE 9




FIGURE 10

7/7

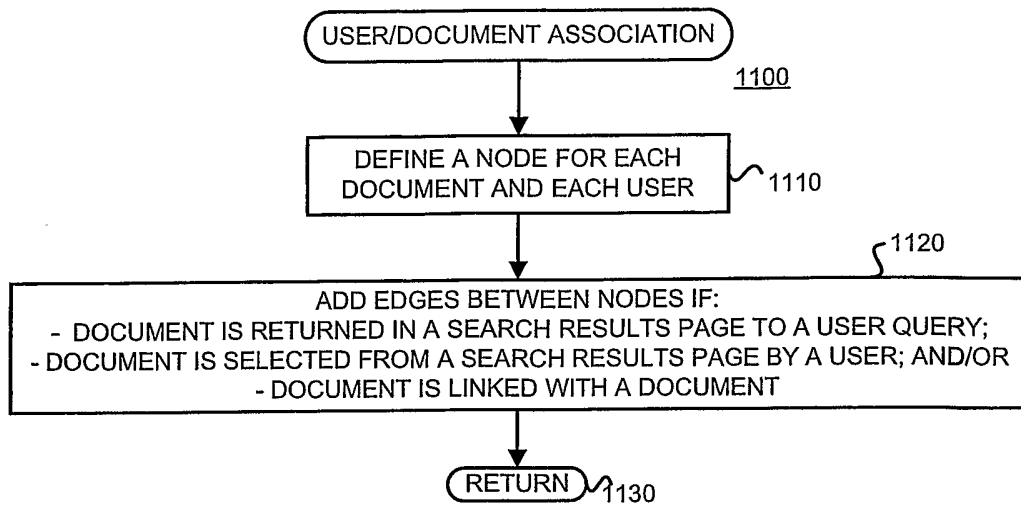


FIGURE 11

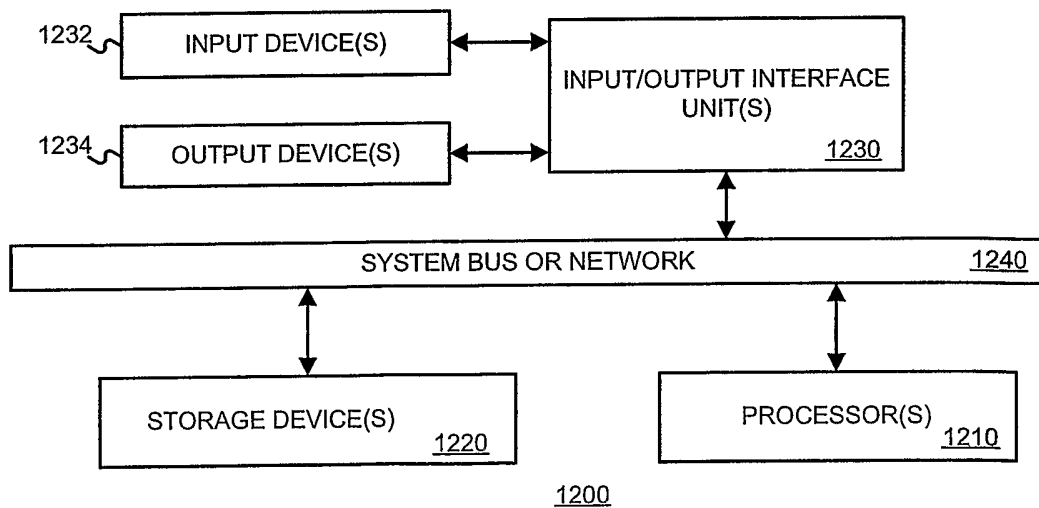


FIGURE 12